

IN THE UNITED STATES DISTRICT COURT
FOR THE EASTERN DISTRICT OF TEXAS
MARSHALL DIVISION

FILED - CLERK
U.S. DISTRICT COURT
SEP 16 PM 3:02
TEXAS-EASTERN
BY _____

REMBRANDT TECHNOLOGIES, LP

Plaintiff,

v.

COMCAST CORPORATION; COMCAST
CABLE COMMUNICATIONS, LLC; and
COMCAST OF PLANO, LP

Defendants.

Case No.

2:05cv443

JURY TRIAL DEMANDED

COMPLAINT FOR PATENT INFRINGEMENT

For its complaint plaintiff Rembrandt Technologies, LP ("Rembrandt"), by and through the undersigned attorneys, alleges as follows:

THE PARTIES

1. Plaintiff Rembrandt is a limited partnership organized under the laws of the state of New Jersey with its principal place of business at 401 City Avenue, Suite 528, Bala Cynwyd, PA 19004.

2. Defendant Comcast Corporation is a corporation organized under the laws of the state of Pennsylvania with its principal place of business at 1500 Market Street, Philadelphia, PA 19102. Defendant Comcast Cable Communications, LLC is an LLC organized under the laws of Delaware with its principal place of business at 1500 Market Street, Philadelphia, PA 19103. Defendant Comcast of Plano, LP is a limited partnership organized under the laws of Delaware. (Defendants are collectively referred to herein as "Comcast.")

JURISDICTION AND VENUE

3. This is an action for patent infringement, arising under the patent laws of the United States, 35 U.S.C. §§ 1, *et seq.*.

4. Subject matter jurisdiction is proper in this Court under 28 U.S.C. §§ 1331 and 1338(a).

5. Because Comcast has committed acts of patent infringement in this judicial district, and is otherwise present or doing business in this judicial district, this Court has personal jurisdiction over Comcast.

6. Because Comcast has committed acts of patent infringement in this judicial district, and because Comcast provides services and does business in this judicial district and otherwise has minimum contacts with this judicial district, venue is proper in this judicial district under 28 U.S.C. §§ 1391(b), (c), and 1400(b).

COUNT I – INFRINGEMENT OF U.S. PATENT NO. 5,243,627

7. Rembrandt realleges and incorporates herein by reference the allegations stated in paragraphs 1-6 of this Complaint.

8. Rembrandt is the owner of all right, title and interest, including the right to sue, enforce and recover damages for all infringements, in U.S. Patent No. 5,243,627, entitled “Signal Point Interleaving Technique” (“the ‘627 patent.”)

9. The ‘627 patent was duly and legally issued by the United States Patent and Trademark Office on September 7, 1993.

10. Comcast is the operator of cable systems throughout the United States.

11. Comcast has directly or indirectly infringed, and is continuing to directly or indirectly infringe, the '627 patent by practicing or causing others to practice (by inducement and contributorily) the inventions claimed in the '627 patent, in this district and throughout the United States. For example, Comcast has infringed and continues to infringe the '627 patent by its receipt and retransmission over its cable television systems of digital terrestrial broadcast signals that comply with the ATSC Digital Television Standard.

12. Upon information and belief, Comcast will continue to infringe the '627 patent unless enjoined by this Court. Upon information and belief, such infringement has been, and will continue to be, willful, making this an exceptional case and entitling Rembrandt to increased damages and reasonable attorneys' fees pursuant to 35 U.S.C. §§ 284 and 285.

COUNT II – INFRINGEMENT OF U.S. PATENT NO. 5,852,631

13. Rembrandt realleges and incorporates herein by reference the allegations stated in paragraphs 1-12 of this Complaint.

14. Rembrandt is the owner of all right, title and interest, including the right to sue, enforce and recover damages for all infringements, in U.S. Patent No. 5,852,631, entitled "System and Method for Establishing Link Layer Parameters Based on Physical Layer Modulation" ("the '631 patent.")

15. The '631 patent was duly and legally issued by the United States Patent and Trademark Office on December 22, 1998.

16. Comcast is the operator of cable systems and provider of Internet service throughout the United States.

17. Comcast has directly or indirectly infringed, and is continuing to directly or indirectly infringe, the '631 patent by practicing or causing others to practice (by inducement and contributorily) the inventions claimed in the '631 patent, in this district and throughout the United States. For example, Comcast has infringed and continues to infringe the '631 patent by its provision of high speed internet service to its customers, including cable subscribers.

18. Upon information and belief, Comcast will continue to infringe the '631 patent unless enjoined by this Court. Upon information and belief, such infringement has been, and will continue to be, willful, making this an exceptional case and entitling Rembrandt to increased damages and reasonable attorneys' fees pursuant to 35 U.S.C. §§ 284 and 285.

COUNT III – INFRINGEMENT OF U.S. PATENT NO. 5,719,858

19. Rembrandt realleges and incorporates herein by reference the allegations stated in paragraphs 1-18 of this Complaint.

20. Rembrandt is the owner of all right, title and interest, including the right to sue, enforce and recover damages for all infringements, in U.S. Patent No. 5,719,858, entitled "Time-Division Multiple-Access Method for Packet Transmission on Shared Synchronous Serial Busses" ("the '858 patent.")

21. The '858 patent was duly and legally issued by the United States Patent and Trademark Office on February 17, 1998.

22. Comcast is the operator of cable systems and provider of Internet service throughout the United States.

23. Comcast has directly or indirectly infringed, and is continuing to directly or indirectly infringe, the '858 patent by practicing or causing others to practice (by inducement and contributorily) the inventions claimed in the '858 patent, in this district and throughout the United States. For example, Comcast has infringed and continues to infringe the '858 patent by its provision of high speed internet services, including such services as Voice over IP (VoIP) services, to its customers, including cable subscribers.

24. Upon information and belief, Comcast will continue to infringe the '858 patent unless enjoined by this Court. Upon information and belief, such infringement has been, and will continue to be, willful, making this an exceptional case and entitling Rembrandt to increased damages and reasonable attorneys' fees pursuant to 35 U.S.C. §§ 284 and 285.

COUNT IV – INFRINGEMENT OF U.S. PATENT NO. 4,937,819

25. Rembrandt realleges and incorporates herein by reference the allegations stated in paragraphs 1-24 of this Complaint.

26. Rembrandt is the owner of all right, title and interest, including the right to sue, enforce and recover damages for all infringements, in U.S. Patent No. 4,937,819, entitled "Time Orthogonal Mutiple Virtual DCE for Use in Analog and Digital Networks" ("the '819 patent.")

27. The '819 patent was duly and legally issued by the United States Patent and Trademark Office on June 26, 1990.

28. Comcast is the operator of cable systems and provider of Internet service throughout the United States.

29. Comcast has directly or indirectly infringed, and is continuing to directly or indirectly infringe, the '819 patent by practicing or causing others to practice (by inducement and contributorily) the inventions claimed in the '819 patent, in this district and throughout the United States. For example, Comcast has infringed and continues to infringe the '819 patent by its provision of high speed internet services, including services such as Voice over IP (VoIP) services, to its customers, including cable subscribers.

30. Upon information and belief, Comcast will continue to infringe the '819 patent unless enjoined by this Court. Upon information and belief, such infringement has been, and will continue to be, willful, making this an exceptional case and entitling Rembrandt to increased damages and reasonable attorneys' fees pursuant to 35 U.S.C. §§ 284 and 285.

PRAYER FOR RELIEF

WHEREFORE, Rembrandt prays that it have judgment against Comcast for the following:

- (1) A decree that Comcast has infringed the patents-in-suit;
- (2) A permanent injunction enjoining and restraining Comcast and its agents, servants, employees, affiliates, divisions, and subsidiaries, and those in association with it, from making, using, offering to sell, selling, and importing into the United States any product, or using, offering to sell, or selling any service, which falls within the scope of any claim of the patents-in-suit;
- (3) An award of damages;
- (4) An award of increased damages pursuant to 35 U.S.C. § 284;
- (5) An award of all costs of this action, including attorneys' fees and interest; and
- (6) Such other and further relief, at law or in equity, to which Rembrandt is justly entitled.

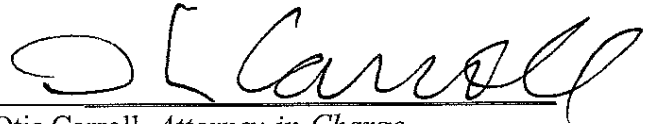
JURY DEMAND

Rembrandt hereby demands a jury trial on all issues appropriately triable by a jury.

Dated: September 16, 2005

Respectfully submitted,

By:



Otis Carroll, *Attorney-in-Charge*
State Bar No. 03895700
Wesley Hill
State Bar No. 24032294
IRELAND, CARROLL & KELLEY, P.C.
6101 S. Broadway, Suite 500
Tyler, Texas 75703
Tel: (903) 561-1600
Fax: (903) 581-1071
Email: fedserv@icklawn.com

Of COUNSEL:

Frank E. Scherkenbach
Lawrence K. Kolodney
FISH & RICHARDSON, P.C.
225 Franklin Street
Boston, Massachusetts 02110
Tel: 617-542-5070
Fax: 617-542-8906

Alan D. Albright
/State Bar # 00973650
FISH & RICHARDSON, P.C.
One Congress Plaza
4th Floor
111 Congress Avenue
Austin, Texas 78701
Tel: 512-391-4930
Fax: 512-391-6837

Timothy Devlin
FISH & RICHARDSON, P.C.
919 N. Market Street, Suite 1100
P.O. Box 1114
Wilmington, Delaware 19899
Tel: 302-652-5070
Fax: 302-652-0607

Franklin Jones, Jr.
State Bar No. 00000055
JONES & JONES, INC., P.C.
201 West Houston Street, Drawer 1249
Marshall, Texas 75671-1249
Tel: (903) 938-4395
Fax: (903) 938-3360

S. Calvin Capshaw, III
State Bar No. 03783900
BROWN MCCARROLL, L.L.P.
1127 Judson Road, Suite 220
P. O. Box 3999
Longview, Texas 75601-5157
Tel: (903) 236-9800
Fax: (903) 236-8787

Robert M. Parker
State Bar No. 15498000
100 E. Ferguson, Suite 1114
Tyler, Texas 75702
Tel: (903) 533-9288
Fax: (903) 533-9687

Attorneys for Plaintiff
REMBRANDT TECHNOLOGIES, LP.



US005243627A

United States Patent [19]

[11] Patent Number: 5,243,627

Betts et al.

[45] Date of Patent: Sep. 7, 1993

[54] SIGNAL POINT INTERLEAVING TECHNIQUE

[56] References Cited

U.S. PATENT DOCUMENTS

[75] Inventors: William L. Betts, St. Petersburg;
Edward S. Zuranski, Largo, both of
Fla.

3,988,677 10/1976 Fletcher et al. 371/45 X
4,677,624 6/1987 Betts et al. 375/39
4,945,549 7/1990 Simon et al. 375/53
5,029,185 7/1991 Wei 375/39 X

[73] Assignee: AT&T Bell Laboratories, Murray
Hill, N.J.

Primary Examiner—Curtis Kuntz
Assistant Examiner—Tesfaldet Bocure
Attorney, Agent, or Firm—Ronald D. Slusky; Gerard A.
deBlasi

[21] Appl. No.: 748,594

[57] ABSTRACT

[22] Filed: Aug. 22, 1991

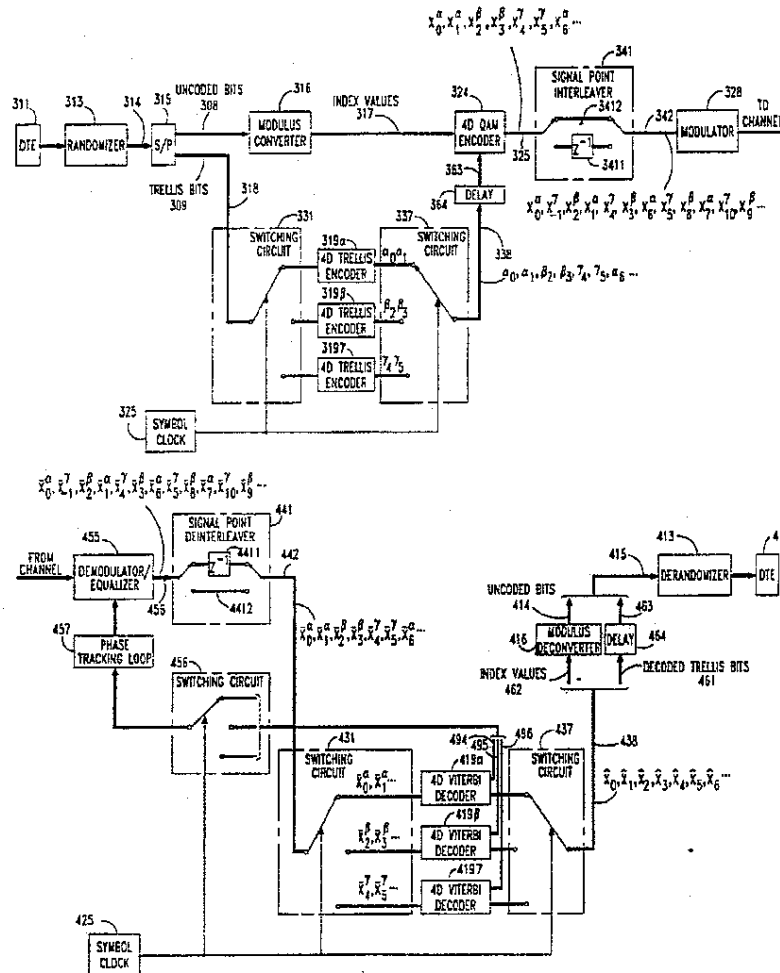
Viterbi decoder performance in a data communication system using 2N-dimensional channel symbols $N > 1$ can be further enhanced by an interleaving technique which uses a distributed trellis encoder in combination with a signal point interleaver.

[51] Int. Cl.⁵ H04L 5/12

[52] U.S. Cl. 375/39; 375/60;
375/99; 371/43

[58] Field of Search 375/39, 58, 60, 99;
371/43, 37.5, 2.1, 45; 341/81

24 Claims, 4 Drawing Sheets



U.S. Patent

Sep. 7, 1993

Sheet 1 of 4

5,243,627

FIG. 1

PRIOR ART

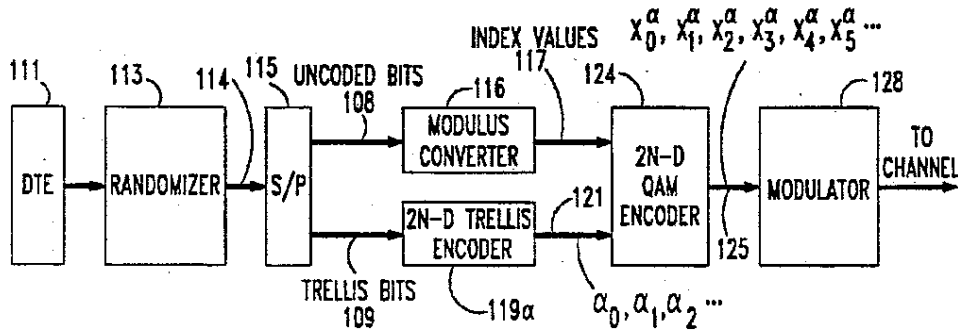



FIG. 2




U.S. Patent

Sep. 7, 1993

Sheet 2 of 4

5,243,627

FIG. 3




U.S. Patent

Sep. 7, 1993

Sheet 3 of 4

5,243,627



U.S. Patent

Sep. 7, 1993

Sheet 4 of 4

5,243,627

FIG. 5

		$\overbrace{x_0^\alpha \ x_1^\alpha}^{4D \text{ SYMBOL}}$	$\overbrace{x_2^\alpha \ x_3^\alpha}^{4D \text{ SYMBOL}}$	$\overbrace{x_4^\alpha \ x_5^\alpha}^{4D \text{ SYMBOL}}$	$\overbrace{x_6^\alpha \ x_7^\alpha}^{4D \text{ SYMBOL}}$	$\overbrace{x_8^\alpha \ x_9^\alpha}^{4D \text{ SYMBOL}}$	$\overbrace{x_{10}^\alpha \ \dots}^{4D \text{ SYMBOL}}$
I	NOT INTERLEAVED ONE TRELLIS STAGE	$x_0^\alpha \ x_1^\alpha$	$x_2^\alpha \ x_3^\alpha$	$x_4^\alpha \ x_5^\alpha$	$x_6^\alpha \ x_7^\alpha$	$x_8^\alpha \ x_9^\alpha$	$x_{10}^\alpha \ \dots$
II	NOT INTERLEAVED THREE TRELLIS STAGES	$x_0^\alpha \ x_1^\alpha$	$x_2^\beta \ x_3^\beta$	$x_4^\gamma \ x_5^\gamma$	$x_6^\alpha \ x_7^\alpha$	$x_8^\beta \ x_9^\beta$	$x_{10}^\gamma \ \dots$
III	INTERLEAVED ONE TRELLIS STAGE	$x_0^\alpha \ x_{-1}^\alpha$	$x_2^\alpha \ x_1^\alpha$	$x_4^\alpha \ x_3^\alpha$	$x_6^\alpha \ x_5^\alpha$	$x_8^\alpha \ x_7^\alpha$	$x_{10}^\alpha \ \dots$
IV	INTERLEAVED TWO TRELLIS STAGES	$x_0^\alpha \ x_{-1}^\beta$	$x_2^\beta \ x_1^\alpha$	$x_4^\alpha \ x_3^\beta$	$x_6^\beta \ x_5^\alpha$	$x_8^\alpha \ x_7^\beta$	$x_{10}^\beta \ \dots$
V	INTERLEAVED THREE TRELLIS STAGES	$x_0^\alpha \ x_{-1}^\gamma$	$x_2^\beta \ x_1^\alpha$	$x_4^\gamma \ x_3^\beta$	$x_6^\alpha \ x_5^\gamma$	$x_8^\beta \ x_7^\alpha$	$x_{10}^\gamma \ \dots$

FIG. 6





FIG. 7



5,243,627

1

SIGNAL POINT INTERLEAVING TECHNIQUE

BACKGROUND OF THE INVENTION

The present invention relates to the transmission of digital data over band-limited channels.

Over the years, the requirements of modern-day digital data transmission over band-limited channels—such as voiceband telephone channels—have resulted in a push for higher and higher bit rates. This push has led to the development and introduction of such innovations as adaptive equalization, multi-dimensional signal constellations, echo cancellation (for two-wire applications), and trellis coding. Today, the data rates achieved using these and other techniques are beginning to approach the theoretical limits of the channel.

It has been found that various channel impairments, whose effects on the achievable bit rate were relatively minor compared to, say, additive white Gaussian noise and linear distortion, have now become of greater concern. These include such impairments as nonlinear distortion and residual (i.e., uncompensated-for) phase jitter. Such impairments are particularly irksome in systems which use trellis coding. Indeed, it has been found that the theoretical improvement in Gaussian noise immunity promised by at least some trellis codes is not realized in real-world applications where these impairments are manifest. The principal reason this is so appears to be that the noise components introduced into the received signal samples are such as to worsen the effectiveness of the Viterbi decoder used in the receiver to recover the transmitted data.

U.S. Pat. No. 4,677,625, issued Jun. 30, 1987 to Betts et al, teaches a method and arrangement in which, through the use of a distributed trellis encoder/Viterbi decoder, the effects of many of these impairments can be reduced. The invention in the Betts et al patent recognizes that a part of the reason that the performance of the Viterbi decoder is degraded by these impairments is the fact that the noise components of channel symbols which closely follow one another in the transmission channel are highly correlated for many types of impairments. And it is that correlation which worsens the effect that these impairments have on the Viterbi decoder. Among the impairments whose noise is correlated in this way are impulse noise, phase "hits" and gain "hits." All of these typically extend over a number of adjacent channel symbols in the channel, and thus all result in channel symbol noise components which are highly correlated. The well-known noise enhancement characteristics of linear equalizers also induce correlated noise in adjacent channel symbols, as does uncompensated-for phase jitter. Also, the occurrence of one of the relatively high power points of the signal constellation can, in pulse code modulation (PCM) systems, for example, give rise to noise on adjacent channel symbols which, again, is correlated.

The Betts et al patent addresses this issue by distributing the outgoing data to a plurality of trellis encoders in round-robin fashion and interleaving the trellis encoder outputs on the transmission channel. In the receiver, the stream of received interleaved channel symbols is correspondingly distributed to a plurality of trellis decoders. Since the successive pairs of channel symbols applied to a particular trellis decoder are separated from one another as they traverse the channel, the correlation of the noise components of these channel symbol

2

pairs is reduced from what it would have otherwise been.

SUMMARY OF THE INVENTION

In accordance with the present invention, it has been realized that the Viterbi decoder performance in a data communication system using $2N$ -dimensional channel symbols can be further enhanced by an interleaving technique which uses, in combination, a) the aforementioned distributed trellis encoder/Viterbi decoder technique and b) a signal point interleaving technique which causes the constituent signal points of the channel symbols to be non-adjacent as they traverse the channel.

In preferred embodiments of the invention, the interleaving is carried out in such a way that every N^{th} signal point in the signal point stream traversing the channel is the N^{th} signal point of a respective one of the channel symbols. This criterion enhances the accuracy with which the phase tracking loop in the receiver performs its function.

Also in preferred embodiments, we have found that the use of three parallel trellis encoders in conjunction with a signal point interleaving regime in which the signal points of each channel symbol are separated from one another by three signaling intervals (bauds) provides an optimum or near-optimum tradeoff between signal point/channel symbol separation and the decoding delay that is caused by the interleaving.

BRIEF DESCRIPTION OF THE DRAWING

In the drawing,

FIG. 1 is a block diagram of the transmitter section of a prior art modem;

FIG. 2 is a signal constellation used by the transmitter of FIG. 1;

FIG. 3 is a block diagram of the transmitter section of a modem employing four-dimensional channel symbols and embodying the principles of the invention;

FIG. 4 is a block diagram of the receiver section of a modem embodying the principles of the invention which processes the received four-dimensional channel symbols generated by the transmitter of FIG. 3;

FIG. 5 is a signal point timing/sequencing chart helpful in explaining the principles of the present invention;

FIG. 6 is a signal point interleaver which can be used in the transmitter of FIG. 3 to interleave the signal points of eight-dimensional channel symbols; and

FIG. 7 is a signal point deinterleaver which can be used in the receiver of FIG. 4 to deinterleave the signal points of eight-dimensional channel symbols.

DETAILED DESCRIPTION

FIG. 1 depicts the transmitter section of a prior art modem employing a $2N$ -dimensional signaling scheme, $N \geq 1$. The modem receives input information in the form of a serial bit stream from data terminal equipment (DTE) 111—illustratively a host computer. That bit stream is then scrambled, or randomized, by randomizer 113 whose output bits are provided in serial form to serial-to-parallel (S/P) converter 115.

Serial-to-parallel converter 115, in turn, provides, during each of a succession of symbol intervals (comprised of N baud intervals), some predetermined number of parallel bits on lead 109 and some number of parallel bits on lead 108. (It will be appreciated that whenever bits are provided in parallel in the modem, separate leads are required to carry each of the bits.) The bits on lead 109 are applied to trellis encoder 119a,

5,243,627

3

and are referred to as the "trellis bits." The bits on lead 108 are applied to modulus converter 116, and are referred to as the "uncoded bits."

To better understand how trellis encoder 119 α and modulus converter 116 work, reference is made to FIG. 2, which shows the two-dimensional signal constellation that forms the basis of the 2N-dimensional signaling scheme illustratively used by the modem. This constellation is comprised of 32 signal points, which are divided into four subsets, A through D, each comprised of eight signal points. The eight points of subset A are explicitly labeled as A₀ through A₇. It may be noted that subsets C, B and D can be arrived at by clockwise rotation of subset A by 90, 180 and 270 degrees, respectively. (Conventional differential encoding circuitry within trellis encoder 119 α exploits this symmetry.) For reference, a single signal point of each of those subsets is also shown on FIG. 2.

Consider, first, the case of N=1, i.e., a two-dimensional signaling scheme. In this case, one trellis bit on lead 109 would be expanded to two bits by trellis encoder 119 α on lead 121. The four possible values of those three bits 00, 01, 10, and 11 identify subsets A, B, C and D, respectively. The successive 2-bit words on lead 121 are represented as α_n , $n=0,1,2, \dots$, where n is an index that advances at the baud rate. At the same time, three parallel bits would be provided on lead 108. These are converted by modulus converter 116 into an index having a value within the range (decimal) 0 to 7. The index value, represented in binary form on lead 117, selects a particular signal point from the subset identified on lead 121. Thus if lead 121 carries the two bits 00 while lead 117 carries the three bits 001, then signal point A₁ of the FIG. 2 constellation has been selected. The words on leads 117 and 121 are applied to QAM encoder 124 which generates, on lead 125, values representing the I (in-phase) and Q (quadrature-phase) components of signal point A₁. The signal point generated on lead 125 in the nth baud interval is denoted X_n α , which is passed on to modulator 128 to generate a pass-band line signal which is applied to the communication channel. The superscript, α , indicates that the trellis encoder that was used to identify the subset for any particular signal point was trellis encoder 119 α . That is, of course, a trivial notation as far as FIG. 1 goes inasmuch as trellis encoder 119 α is the only trellis encoder in the modem. However, it is useful to introduce this notation because more than one trellis encoder stage is used in preferred embodiments of modems incorporating the principles of the present invention as shown in later FIGS.

In the case of N>1, the operation is similar. Now, however, the words on lead 109 are used by trellis encoder 119 α to sequentially identify on lead 121N subsets, while the words on lead 108 are used to generate N corresponding index values on lead 117. The N signal points identified in this way are the component signal points of a 2N-dimensional channel symbol, the first such symbol being comprised of the signal points X₀ α , \dots , X_(N-1) α . For example, a modem in which the transmitter of FIG. 1 could be used may be a 14,400 bit per second modem using four-dimensional coding (i.e., N=2) and a baud rate of 3200. In this case, nine bits from S/P converter 115 are used for each four-dimensional symbol. Specifically, three parallel bits on lead 109 are expanded into four bits on lead 121 to identify a pair of subsets while six bits on lead 108 are used to select particular signal points from those two subsets.

4

Those two signal points are thereupon communicated over the channel by QAM encoder 124 and modulator 128 as described above.

Note that, implementationally, the 2N-dimensional channel symbol is generated by having the trellis encoder identify, interdependently, N subsets of the two-dimensional constellation of FIG. 2, then select a two-dimensional signal point from each of the subsets thus identified. The concatenation of the N two-dimensional signal points thus selected is the desired 2N-dimensional channel symbol. This process, however, can be understood as involving the direct selection of a 2N-dimensional channel symbol. Viewed in this context, the set of all possible combinations of N of the two-dimensional subsets identified by N successive trellis encoder outputs can be understood to be a set of 2N-dimensional subsets of a 2N-dimensional constellation, the latter being comprised of all possible combinations of N of the signal points of the two-dimensional constellation. A succession of N outputs from the trellis encoder identifies a particular one of the 2N-dimensional subsets and a succession of N outputs from the modulus converter selects a particular 2N-dimensional signal point from the identified 2N-dimensional subset.

Modulus converter 116 is illustratively of the type disclosed in co-pending, commonly-assigned U.S. patent application Ser. No. 588,658 filed Sep. 26, 1990 and allowed on May 21, 1991, hereby incorporated by reference. Modulus converter 116 provides the modem with the ability to support data transmission at various different bit rates. Assume, for example, that the rate at which bits are provided by DTE 111 decreases. The serial-to-parallel converter will continue to provide its outputs on leads 108 and 109 at the same baud rate as before. However, the upper limit of the range of index values that are provided by modulus converter 116 on lead 117 will be reduced, so that, effectively, each of the four subsets A through D, instead of having eight signal points, will have some smaller number. Conversely if the rate at which bits are provided by DTE 111 should increase over that originally assumed, the upper limit of the range of index values, and thus the number of parallel bits, that appear on lead 117 will be increased beyond eight and the constellation itself will be expanded to accommodate the larger number of signal points thus being selected. As an alternative to using a modulus converter, fractional bit rates can be supported using, for example, the technique disclosed in L. Wei, "Trellis-Coded Modulation with Multidimensional Constellations," *IEEE Trans. on Communication Theory*, Vol. IT-33, No. 4, July 1987, pp. 483-501.

Turning now to FIG. 3, the transmitter portion of a modem embodying the principles of the invention is shown. This embodiment illustratively uses the aforementioned four-dimensional, i.e., N=2, signaling scheme. Many of the components are similar to those shown in FIG. 1. Thus, in particular, the transmitter of FIG. 3—which receives its input information in the form of a stream of input bits from DTE 311—includes randomizer 313, which supplies its output, on lead 314, to S/P converter 315. The latter outputs uncoded bits to modulus converter 316. The transmitter further includes four-dimensional QAM encoder 324 and modulator 328. The trellis bits, on lead 309, are provided not to a standard single trellis encoder, but to a distributed trellis encoder comprised of three trellis encoder stages: trellis encoder stage 319 α , trellis encoder stage 319 β , and trellis encoder stage 319 γ .

5,243,627

5

Such a distributed trellis encoder, which is described in the aforementioned Betts et al patent, generates a plurality of streams of trellis encoded channel symbols in response to respective portions of the input information. Specifically, a three-bit word on lead 309 is supplied to trellis encoder stage 319 α . The next three-bit word on lead 309 is supplied to trellis encoder stage 319 β . The next three-bit word is supplied to trellis encoder stage 319 γ , and then back to trellis encoder stage 319 α . This distribution of the trellis bits to the various trellis encoder stages is performed by switching circuit 331 operating under the control of symbol clock 325. The initial data word outputs of the trellis encoders are subset identifiers α_0 and α_1 for encoder stage 319 α , β_2 and β_3 for encoder stage 319 β , and γ_4 and γ_5 for encoder stage 319 γ , followed by α_6 and α_7 for encoder stage 319 α , and so forth. These are supplied to four-dimensional QAM encoder 324 by switching circuit 337—also operating under the control of symbol clock 325—on lead 338 through a one-symbol delay 364 and lead 363, in order to compensate for a one-symbol delay caused by modulus converter 316. Thus, the stream of subset identifiers on lead 338 is $\alpha_0, \alpha_1, \beta_2, \beta_3, \gamma_4, \gamma_5, \alpha_6, \dots$. Using the notation introduced above, then, the output of encoder 324 on lead 325 is the stream of signal points $X_0^\alpha, X_1^\alpha, X_2^\beta, X_3^\beta, X_4^\gamma, X_5^\gamma, X_6^\alpha, \dots$, which is comprised of three interleaved streams of trellis encoded channel symbols, these streams being $X_0^\alpha, X_1^\alpha, X_6^\alpha, X_7^\alpha, X_{12}^\alpha, \dots$; $X_2^\beta, X_3^\beta, X_8^\beta, X_9^\beta, X_{14}^\beta, \dots$; and $X_4^\gamma, X_5^\gamma, X_{10}^\gamma, X_{11}^\gamma, X_{16}^\gamma, \dots$. These, in turn, are supplied, in accordance with the invention, to signal point interleaver 341 which applies alternate ones of the signal points applied thereto to lead 3412—which signal points appear immediately at the interleaver output on lead 342—and to one-symbol (Z^{-1}) delay element 3411, which appear on lead 342 after being delayed therein by one symbol interval. The resulting interleaved stream of trellis encoded signal points is $X_0^\alpha, X_{-1}^\gamma, X_2^\beta, X_1^\alpha, X_4^\gamma, X_3^\beta, X_6^\alpha, X_5^\gamma, X_8^\beta, X_7^\alpha, X_{10}^\gamma, X_9^\beta, \dots$ (the signal point X_{-1}^γ being, of course, the signal point applied to interleaver 341 just ahead of signal point X_0^α).

A discussion and explanation of how the interleaving just described is advantageous is set forth hereinbelow. In order to fully set the stage for that explanation, however, it will be first useful to consider the receiver section of a modem which receives the interleaved signal point stream.

Thus referring to FIG. 4, the line signal transmitted by the transmitter of FIG. 3 is received from the channel and applied to demodulator/equalizer 455 which, in conventional fashion—including an input from phase tracking loop 457—generates a stream of outputs on lead 456 representing the demodulator/equalizer's best approximation of the values of the I and Q components of the signal points of the transmitted interleaved signal point stream. These outputs are referred to herein as the "received signal points." (Due to distortion and other channel impairments that the demodulator/equalizer is not able to compensate for, the I and Q components of the received signal points, instead of having exact integer values, can have any value. Thus a transmitted signal point having coordinates (3, -5) may be output by the demodulator/equalizer as the received signal point (2.945, -5.001).) The stream of received signal points on lead 456 is denoted $\bar{X}_0^\alpha, \bar{X}_{-1}^\gamma, \bar{X}_2^\beta, \bar{X}_1^\alpha, \bar{X}_4^\gamma, \bar{X}_3^\beta, \bar{X}_6^\alpha, \bar{X}_5^\gamma, \bar{X}_8^\beta, \bar{X}_7^\alpha, \bar{X}_{10}^\gamma, \bar{X}_9^\beta, \dots$

The successive received signal points are deinterleaved in signal point deinterleaver 441, which provides

6

the opposite function to interleaver 341 in the transmitter. The output of deinterleaver 441 on lead 442 is thus $\bar{X}_0^\alpha, \bar{X}_1^\alpha, \bar{X}_2^\beta, \bar{X}_3^\beta, \bar{X}_4^\gamma, \bar{X}_5^\gamma, \bar{X}_6^\alpha, \dots$, etc. (Although not explicitly shown in the drawing, the same well-known techniques used in modems of this general kind to identify within the stream of received signal points the boundaries between successive symbols is used to synchronize the operation of signal point deinterleaver 441 to ensure that received signal points $\bar{X}_0^\alpha, \bar{X}_2^\beta, \bar{X}_4^\gamma, \dots$ are applied to delay element 4411 while received signal points $\bar{X}_1^\alpha, \bar{X}_3^\beta, \bar{X}_5^\gamma, \dots$ are applied to lead 4412.)

The received signal points on lead 442 are then distributed by switching circuit 431 under the control of symbol clock 425 to a distributed Viterbi decoder comprised of 4D Viterbi decoder stages 419 α , 419 β and 419 γ . Specifically, received signal points \bar{X}_0^α and \bar{X}_1^α are applied to decoder stage 419 α ; received signal points \bar{X}_2^β and \bar{X}_3^β are applied to decoder stage 419 β ; and received signal points \bar{X}_4^γ and \bar{X}_5^γ are applied to decoder stage 419 γ . The outputs of the three decoder stages are then combined into a serial stream on lead 438 by switching circuit 437, also operating under the control of symbol clock 425. Those outputs, representing decisions as to the values of the transmitted signal points, are denoted $\hat{X}_0, \hat{X}_1, \hat{X}_2, \hat{X}_3, \hat{X}_4, \hat{X}_5, \hat{X}_6, \dots$, the α, β and γ superscripts no longer being needed.

In conventional fashion, the bits that represent each of the decisions on lead 438 can be divided into bits that represent a) the trellis bits that appeared on transmitter lead 309 and b) the index values that appeared on transmitter lead 317. Those two groups of bits are provided in the receiver on leads 461 and 462, respectively. The latter group of bits are deconverted by modulus deconverter 416 (also disclosed in the aforementioned '658 patent application) back to uncoded bit values on lead 414. The operation of the modulus deconverter imparts a one-symbol delay to the bits on lead 414. Accordingly, the bits on lead 461 are caused to be delayed by one symbol by delay element 464. The resulting combined bits on lead 415 thus represent the stream of bits that appeared at the output of randomizer 313 in the transmitter. These are derandomized in the receiver by derandomizer 413 and the resulting derandomized bit stream is applied to DTE 411 which may be, for example, a computer terminal.

Referring to FIG. 5, one can see the improvement that is achieved by the present invention.

Line I shows the stream of output signal points generated and launched into the channel using one stage of trellis encoding and no signal point interleaving. This is, of course, the prior art arrangement shown in FIG. 1. Line II shows the effect of providing a three-stage distributed trellis encoder but still no signal point interleaving. This is the arrangement shown in the aforementioned Betts et al patent. Note that the signal points of each channel symbol operated on by a particular trellis encoder stage are adjacent in the output signal point stream. For example, the second signal point of the symbol $X_0^\alpha X_1^\alpha$ —namely signal point X_1^α —is separated by five baud intervals from the first (closer) signal point of the symbol $X_6^\alpha X_7^\alpha$ —namely signal point X_6^α . As noted earlier, such separation is advantageous because the channel symbols which are processed one after the other in a particular Viterbi decoder stage have noise components which are not highly correlated.

Note, however, that the individual signal points of each channel symbol, e.g., X_0^α and X_1^α , are adjacent to

5,243,627

7

one another as they pass through the channel; and since all the signal points of a channel symbol must be processed serially in the same Viterbi decoder stage, this means that the Viterbi decoder must process adjacent signal points that have highly correlated noise components.

It is to this end that signal point interleaver 341 is included within the transmitter in accordance with the invention. Firstly, it may be noted from Line III that using the signal point interleaver without the distributed trellis encoder—an arrangement not depicted in the drawing—will, advantageously, cause the signal points from the same channel symbol to be non-adjacent. Moreover, there is further advantage in that a pair of channel symbols processed serially by Viterbi decoder stage 419 α traverses the channel separated by five baud intervals rather than three, thereby providing greater decorrelation of the noise components thereof. Compare, for example, the span of baud intervals occupied by signal points X_0^α and X_1^α , X_2^α and X_3^α in Line I and the span of baud intervals occupied by the same signal points in Line III. Disadvantageously, however, the use of a single trellis encoding stage brings back the problem that the distributed trellis encoder solves, as described above. Thus, for example, although signal points X_0^α and X_1^α , which are from the same channel symbol, are separated from one another when traversing the channel, we find that, disadvantageously, signal points X_2^α and X_1^α , which are signal points from two different channel symbols which will be processed serially by the Viterbi decoder, traverse the channel adjacent to one another.

Line IV shows that using the signal point interleaver with a two-stage trellis encoder—also an arrangement not depicted in the drawing—provides some improvement. Firstly, it may be noted that, as in Line III, signal points from the same channel symbol remain separated by three baud intervals. Additionally, pairs of channel symbols processed sequentially by a given Viterbi decoder stage—such as the channel symbols comprised of signal points X_0^α and X_1^α , X_2^α and X_3^α —are still non-adjacent and, indeed, are now separated by seven baud intervals, which is even greater than the separation of five baud intervals provided in Line III. Moreover, certain signal points that traverse the channel adjacent to one another and which are from channel symbols which would have been decoded sequentially in the one-trellis-encoding-stage case are, in the two-trellis-encoding-stage case of Line IV, processed by different Viterbi decoding stages. Signal points X_2^β and X_1^α are such a pair of signal points. Note, however, that, disadvantageously, signal points X_1^α and X_4^α traverse the channel serially, and are from channel symbols which are serially processed by the “ α ” Viterbi decoder stage.

Referring, however, to Line V, which depicts the stream of signal points output by the transmitter of FIG. 3, it will be seen that, in accordance with the invention, there is still a non-adjacency—indeed, a separation of at least three baud intervals—between a) the signal points which belong to any particular channel symbol (and which, therefore, are processed serially by a particular Viterbi decoder stage) and b) the signal points which belong to channel symbols which are processed serially by a Viterbi decoder stage. Thus, for example, signal points X_1^α and X_4^α are now processed by different Viterbi decoder stages. Moreover, pairs of channel symbols processed sequentially by a given Viterbi decoder stage—such as the channel symbols comprised of

8

signal points X_0^α and X_1^α , X_6^α and X_7^α —are now separated by none baud intervals.

Using more than three trellis encoder stages in the distributed trellis encoder and/or a signal point interleaver that separates signal points from the same channel symbol by more than three baud intervals would provide even greater separation and could, therefore, potentially provide even greater improvement in Viterbi decoding. However, such improvement comes at a price—that price being increased decoding delay—particularly as the number of trellis encoders is increased beyond three. An engineering trade-off can be made, as suits any particular application.

Moreover, it is desirable for the signal point interleaver to provide a sequence in which every N^{th} signal point in the interleaved signal point stream is the N^{th} signal point of a channel symbol. (The reason this is desirable is described in detail hereinbelow.) In the case of an $N=2$, four-dimensional signaling scheme, this means that every second, that is “every other,” signal point in the interleaved stream is the second signal point of the channel symbol from which it comes. In the case of an $N=4$, eight-dimensional signaling scheme, this means that every fourth signal point in the interleaved stream is the fourth signal point of the channel symbol from which it comes. Indeed, this criterion is in fact satisfied in the embodiment of FIG. 3. Note that each one of signal points X_0^α , X_2^β , X_4^γ , X_6^α , . . . , which appear as every other signal point in the interleaved stream, is the second signal point of one of the four-dimensional channel symbols. Note that not all rearrangements of the signal points will, in fact, satisfy this criterion, such as, if the two signal points of a channel symbol are separated by two, rather than three, baud intervals.

Satisfying the above criterion is advantageous because it enhances the accuracy with which phase tracking loop 457 performs its function. This is so because the arrival of an N^{th} signal point of a given symbol means that all the signal points comprising that channel symbol have arrived. This, in turn, makes it possible to form a decision as to the identity of that channel symbol by using the minimum accumulated path metric in the Viterbi decoder stages. (Those decisions are fed back to the tracking loop by decoder stages 419 α , 419 β 419 γ on leads 494, 495 and 496, respectively, via switching circuit 456.) Without having received all of the signal points of a channel symbol, one cannot take advantage of the accumulated path metric information but, rather, must rely on the so-called raw sliced values, which is less accurate. By having every N^{th} signal point in the interleaved stream be the N^{th} signal point of a channel symbol, we are guaranteed that the time between adjacent such path metric “decisions” supplied to the phase tracking loop is, advantageously, never more than N baud intervals.

The foregoing merely illustrates the principles of the invention. Thus although the illustrative embodiment utilizes a four-dimensional signaling scheme, the invention can be used with signaling schemes of any dimensionality. In the general, $2N$ -dimensional, case each stage of the distributed trellis encoder would provide N two-dimensional subset identifiers to switching circuit 337 before the latter moves on to the next stage. And, of course, each stage of the distributed Viterbi decoder would receive N successive received signal points. The distributed trellis encoder and distributed Viterbi decoder can, however, continue to include three trellis

5,243,627

9

encoders and still maintain, independent of the value of N , a separation of three baud intervals in the channel between signal points that are from channel symbols that are adjacent in the trellis encoder. If a greater separation of such signal points is desired, more stages can be added to the distributed trellis encoder/Viterbi decoder, just as was noted above for the four-dimensional case. However, when dealing with $2N$ -dimensional signaling where $N > 2$, it is necessary to add additional delay elements to the signal point interleaver/deinterleaver in order to maintain a three-baud-interval separation among the signal points from any given channel symbol.

Consider, for example, the case of $N=4$, i.e., an eight-dimensional case. Looking again at FIG. 3, the three (8D) stages of the distributed trellis encoder would generate the three streams of subset identifiers $\alpha_0 \alpha_1 \alpha_2 \alpha_3 \alpha_{12} \dots, \beta_4 \beta_5 \beta_6 \beta_7 \beta_{16} \dots$, and $\gamma_8 \gamma_9 \gamma_{10} \gamma_{11} \gamma_{20} \dots$, respectively. This would lead to the following stream of signal points of eight-dimensional trellis encoded channel symbols at the output of the QAM encoder on lead 325: $X_0^\alpha X_1^\alpha X_2^\alpha X_3^\alpha X_4^\beta X_5^\beta X_6^\beta X_7^\beta X_8^\gamma X_9^\gamma X_{10}^\gamma X_{11}^\gamma X_{12}^\alpha \dots$. Signal point interleaving could be carried out by substituting signal point interleaver 641 of FIG. 6 for interleaver 341. Interleaver 641, in addition to direct connection 6414, includes one-, two-, and three-symbol delay elements 6413, 6412 and 6411, respectively.

The signal points on lead 325, after passing through interleaver 641, would appear on lead 342 in the following order: $X_0^\alpha X_{-3}^\gamma X_{-6}^\beta X_{-9}^\alpha X_4^\beta X_1^\alpha X_{-2}^\gamma X_{-5}^\beta X_8^\gamma X_5^\beta X_2^\alpha X_{-1}^\gamma X_{12}^\alpha X_9^\gamma X_6^\beta X_3^\alpha X_{16}^\beta X_{13}^\alpha X_{10}^\gamma X_7^\beta \dots$ where signal points with negative subscripts are, of course, signal points that arrived before signal point X_0^α and were already stored in the delay elements 6411, 6412 and 6413. Examination of this signal point stream will reveal that there is either a three- or five-baud separation between signal points of channel symbols that are processed sequentially by the same trellis encoder stage, e.g., X_3^α and X_{12}^α ; that adjacent signal points of any one channel symbol, e.g., X_0^α and X_1^α , are separated by five baud intervals; and that the four signal points comprising any particular one channel symbol are separated by fifteen baud intervals.

FIG. 7 shows the structure of a deinterleaver 741 that could be used in the receiver of FIG. 4 in place of deinterleaver 441 in order to restore the signal points of the eight-dimensional channel symbols to their original order. This structure, which is the inverse of interleaver 641, includes delay stages 7411, 7412 and 7413, as well as direct connection 7414.

It will be appreciated that, although various components of the modem transmitter and receiver are disclosed herein for pedagogic clarity as discrete functional elements and indeed—in the case of the various switching circuits—as mechanical elements, those skilled in the art will recognize that the function of any one or more of those elements could be implemented with any appropriate available technology, including one or more appropriately programmed processors, digital signal processing (DSP) chips, etc. For example, multiple trellis encoders and decoders can be realized using a single program routine which, through the mechanism of indirect addressing of multiple arrays within memory, serves to provide the function of each of the multiple devices.

It will thus be appreciated that those skilled in the art will be able to devise numerous arrangements which,

10

although not explicitly shown or described herein, embody the principles of the invention and are within its spirit and scope.

We claim:

1. Apparatus for forming a stream of trellis encoded signal points in response to input information, said apparatus comprising

means for generating a plurality of streams of trellis encoded channel symbols in response to respective portions of said input information, each of said channel symbols being comprised of a plurality of signal points, and

means for interleaving the signal points of said generated channel symbols to form said stream of trellis encoded signal points, said interleaving being carried out in such a way that the signal points of each channel symbol are non-adjacent in said stream of trellis encoded signal points and such that the signal points of adjacent symbols in any one of said channel symbol streams are non-adjacent in said stream of trellis encoded signal points.

2. The apparatus of claim 1 wherein said means for generating generates three of said streams of trellis encoded channel symbols, and wherein said means for interleaving causes there to be interleaved between each of the signal points of each channel symbol at least two signal points from other channel symbols of said streams of trellis encoded channel symbols.

3. The apparatus of claim 1 wherein said channel symbols are $2N$ -dimensional channel symbols, $N > 1$, and wherein said means for interleaving causes every N^{th} signal point in said interleaved signal point stream to be the N^{th} signal point of a respective one of said channel symbols.

4. The apparatus of claim 2 wherein said channel symbols are $2N$ -dimensional channel symbols, $N > 1$, and wherein said means for interleaving causes every N^{th} signal point in said interleaved signal point stream to be the N^{th} signal point of a respective one of said channel symbols.

5. A modem comprising

means for receiving a stream of input bits,

means for dividing said stream of input bits into a stream of uncoded bits and a plurality of streams of trellis bits,

means for independently trellis encoding each of said plurality of streams of trellis bits to generate respective streams of data words each identifying one of a plurality of predetermined subsets of the channel symbols of a predetermined $2N$ -dimensional constellation, N being an integer greater than unity, each of said channel symbols being comprised of a plurality of signal points,

means for selecting an individual channel symbol from each identified subset in response to said stream of uncoded bits to form a stream of channel symbols, and

means for generating a stream of output signal points, said signal point stream being comprised of the signal points of the selected channel symbols, the signal points of said signal point stream being sequenced in such a way that signal points that are either a) part of the same channel symbol, or b) part of channel symbols that are adjacent to one another in said channel symbol stream, are separated in said output stream by at least one other signal point.

6. The apparatus of claim 5 wherein said trellis encoding means includes a plurality of trellis encoder stage

5,243,627

11

means for trellis encoding respective ones of said streams of trellis bits.

7. The apparatus of claim 5 wherein said means for selecting includes means for modulus converting said stream of uncoded bits.

8. The apparatus of claim 5 wherein said channel symbols are 2N-dimensional channel symbols, $N > 1$, and wherein said means for generating causes every N^{th} signal point in said stream of output signal points to be the N^{th} signal point of a respective one of said channel symbols.

9. Receiver apparatus for recovering information from a received stream of trellis encoded signal points, said signal points having been transmitted to said receiver apparatus by transmitter apparatus which generates said signal points by generating a plurality of streams of trellis encoded channel symbols in response to respective portions of said information, each of said channel symbols being comprised of a plurality of signal points, and by interleaving the signal points of said generated channel symbols to form said stream of trellis encoded signal points, said interleaving being carried out in such a way that the signal points of each channel symbol are non-adjacent in said stream of trellis encoded signal points and such that the signal points of adjacent symbols in any one of said channel symbol streams are non-adjacent in said stream of trellis encoded signal points,

said receiver apparatus comprising means for deinterleaving the interleaved signal points to recover said plurality of streams of trellis encoded channel symbols, and a distributed Viterbi decoder for recovering said information from the deinterleaved signal points.

10. The apparatus of claim 9 further comprising a phase tracking loop, and means for adapting the operation of said phase tracking loop in response to minimum accumulated path metrics in said distributed Viterbi decoder.

11. A method for forming a stream of trellis encoded signal points in response to input information, said method comprising the steps of

generating a plurality of streams of trellis encoded channel symbols in response to respective portions of said input information, each of said channel symbols being comprised of a plurality of signal points, and

interleaving the signal points of said generated channel symbols to form said stream of trellis encoded signal points, said interleaving being carried out in such a way that the signal points of each channel symbol are non-adjacent in said stream of trellis encoded signal points and such that the signal points of adjacent symbols in any one of said channel symbol streams are non-adjacent in said stream of trellis encoded signal points.

12. The method of claim 11 wherein said generating step generates three of said streams of trellis encoded channel symbols, and wherein said interleaving step causes there to be interleaved between each of the signal points of each channel symbol at least two signal points from other channel symbols of said streams of trellis encoded channel symbols.

13. The method of claim 11 wherein said channel symbols are 2N-dimensional channel symbols, $N > 1$, and wherein said interleaving step causes every N^{th} signal point in said interleaved signal point stream to be

12

the N^{th} signal point of a respective one of said channel symbols.

14. The method of claim 12 wherein said channel symbols are 2N-dimensional channel symbols, $N > 1$, and wherein said interleaving step causes every N^{th} signal point in said interleaved signal point stream to be the N^{th} signal point of a respective one of said channel symbols.

15. A method for use in a modem, said method comprising the steps of

receiving a stream of input bits, dividing said stream of input bits into a stream of uncoded bits and a plurality of streams of trellis bits,

independently trellis encoding each of said plurality of streams of trellis bits to generate respective streams of data words each identifying one of a plurality of predetermined subsets of the channel symbols of a predetermined 2N-dimensional constellation, N being an integer greater than unity, each of said channel symbols being comprised of a plurality of signal points,

selecting an individual channel symbol from each identified subset in response to said stream of uncoded bits to form a stream of channel symbols, and

generating a stream of output signal points, said signal point stream being comprised of the signal points of the selected channel symbols, the signal points of said signal point stream being sequenced in such a way that signal points that are either a) part of the same channel symbol, or b) part of channel symbols that are adjacent to one another in said channel symbol stream, are separated in said output stream by at least one other signal point.

16. The method of claim 15 wherein in said trellis encoding step a plurality of trellis encoder stages trellis encode respective ones of said streams of trellis bits.

17. The method of claim 15 wherein said selecting step includes the step of modulus converting said stream of uncoded bits.

18. The method of claim 15 wherein said channel symbols are 2N-dimensional channel symbols, $N > 1$, and wherein said generating step causes every N^{th} signal point in said stream of output signal points to be the N^{th} signal point of a respective one of said channel symbols.

19. A method for use in a receiver to recover information from a received stream of trellis encoded signal points, said signal points having been transmitted to said receiver apparatus by a method which includes the steps of

generating a plurality of streams of trellis encoded channel symbols in response to respective portions of said information, each of said channel symbols being comprised of a plurality of signal points, and interleaving the signal points of said generated channel symbols to form said stream of trellis encoded signal points, said interleaving being carried out in such a way that the signal points of each channel symbol are non-adjacent in said stream of trellis encoded signal points and such that the signal points of adjacent symbols in any one of said channel symbol streams are non-adjacent in said stream of trellis encoded signal points,

said method comprising the steps of deinterleaving the interleaved signal points to recover said plurality of streams of trellis encoded channel symbols, and

5,243,627

13

using a distributed Viterbi decoder to recover said information from the deinterleaved signal points.

20. The method of claim 19 wherein said receiver includes a phase tracking loop and wherein said method comprises the further step of adapting the operation of said phase tracking loop in response to minimum accumulated path metrics in said distributed Viterbi decoder.

21. Data communication apparatus comprising means for receiving input information,

means for generating a plurality of streams of trellis encoded channel symbols in response to respective portions of said input information, each of said channel symbols being comprised of a plurality of signal points,

means for interleaving the signal points of said generated channel symbols to form a stream of trellis encoded signal points, said interleaving being carried out in such a way that the signal points of each channel symbol are non-adjacent in said stream of trellis encoded signal points and such that the signal points of adjacent symbols in any one of said channel symbol streams are non-adjacent in said stream of trellis encoded signal points,

means for applying the stream of trellis encoded signal points to a transmission channel,

14

means for receiving the stream of trellis encoded signal points from the channel,

means for deinterleaving the interleaved signal points to recover said plurality of streams of trellis encoded channel symbols, and

a distributed Viterbi decoder for recovering said information from the deinterleaved signal points.

22. The apparatus of claim 21 wherein said means for generating generates three of said streams of trellis encoded channel symbols, and wherein said means for interleaving causes there to be interleaved between each of the signal points of each channel symbol at least two signal points from other channel symbols of said streams of trellis encoded channel symbols.

23. The apparatus of claim 21 wherein said channel symbols are $2N$ -dimensional channel symbols, $N > 1$, and wherein said means for interleaving causes every N^{th} signal point in said interleaved signal point stream to be the N^{th} signal point of a respective one of said channel symbols.

24. The apparatus of claim 22 wherein said channel symbols are $2N$ -dimensional channel symbols, $N > 1$, and wherein said means for interleaving causes every N^{th} signal point in said interleaved signal point stream to be the N^{th} signal point of a respective one of said channel symbols.

* * * * *

30

35

40

45

50

55

60

65



US005852631A

United States Patent [19]

[11] Patent Number: 5,852,631

Scott

[45] Date of Patent: *Dec. 22, 1998

[54] SYSTEM AND METHOD FOR
ESTABLISHING LINK LAYER
PARAMETERS BASED ON PHYSICAL
LAYER MODULATION

5,425,080 6/1995 Abbie 379/91.01
5,491,720 2/1996 Davis et al. 375/222
5,710,761 1/1998 Scott 370/252

FOREIGN PATENT DOCUMENTS

[75] Inventor: Robert Earl Scott, Indian Rocks
Beach, Fla.

0 418 165 3/1991 European Pat. Off. 375/222

[73] Assignee: Paradyne Corporation, Largo, Fla.

Primary Examiner—Stephen Chin
Assistant Examiner—Jeffrey W. Gluck
Attorney, Agent, or Firm—Thomas, Kayden, Horstemeier
& Risley, L.L.P.

[*] Notice: The term of this patent shall not extend
beyond the expiration date of Pat. No.
5,710,761.

[57] ABSTRACT

[21] Appl. No.: 780,762

A system and method for establishing a link layer connection between a calling modem having a plurality of possible first physical layer modulations and one or more possible link layer connections and an answering modem having a plurality of possible second physical layer modulations and one or more possible link layer connections comprising the steps of establishing a physical layer connection between the calling and the answering modems, wherein the physical layer connection is based on a negotiated physical layer modulation chosen from the first and second physical layer modulations, and establishing link layer connection based upon said negotiated physical layer modulation. The link layer connection includes parameters that are preset to default values based upon the negotiated physical layer connection. Thus, the modems are able to avoid the link layer negotiation, thereby providing a faster and more robust connection.

[22] Filed: Jan. 8, 1997

Related U.S. Application Data

[60] Provisional application No. 60/026,970, Sep. 20, 1996, and
provisional application No. 60/022,474, Jun. 21, 1996.

[51] Int. Cl.⁶ H04L 29/10

[52] U.S. Cl. 375/222; 379/93.32

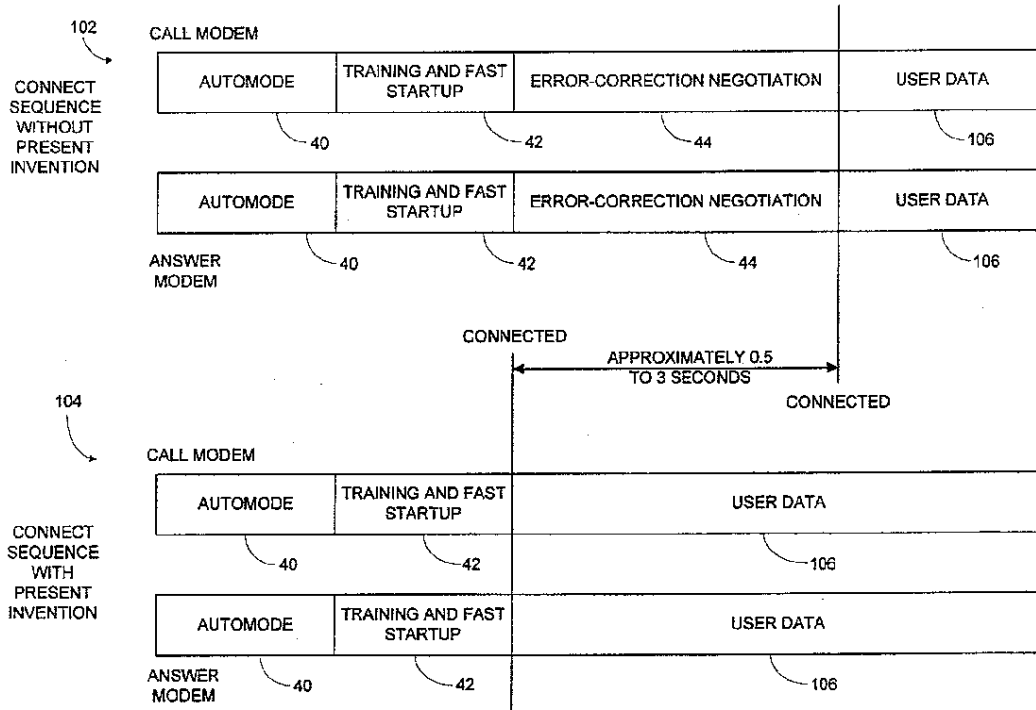
[58] Field of Search 375/222, 223;
379/93.29, 93.31, 93.32, 93.33, 120; 370/252

[56] References Cited

U.S. PATENT DOCUMENTS

4,905,282 2/1990 McGlynn et al. 380/48
4,931,250 6/1990 Greszczuk 375/222
5,317,594 5/1994 Goldstein 375/222

10 Claims, 8 Drawing Sheets



U.S. Patent

Dec. 22, 1998

Sheet 1 of 8

5,852,631




FIG. 1

U.S. Patent

Dec. 22, 1998

Sheet 2 of 8

5,852,631




FIG. 2




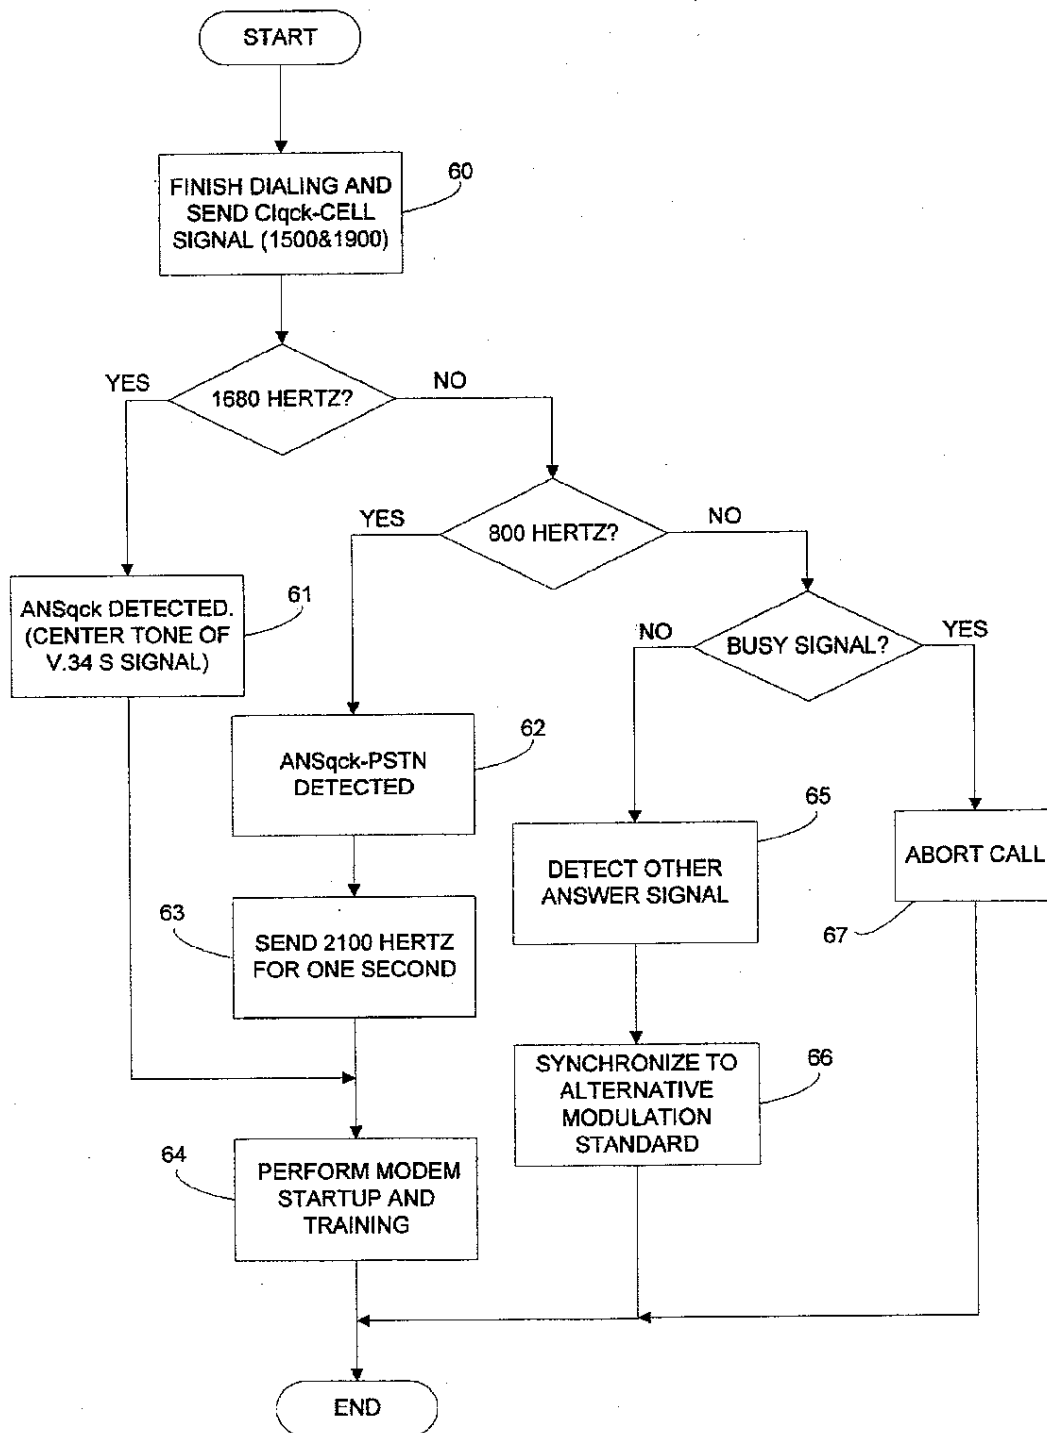
FIG. 3

U.S. Patent

Dec. 22, 1998

Sheet 3 of 8

5,852,631



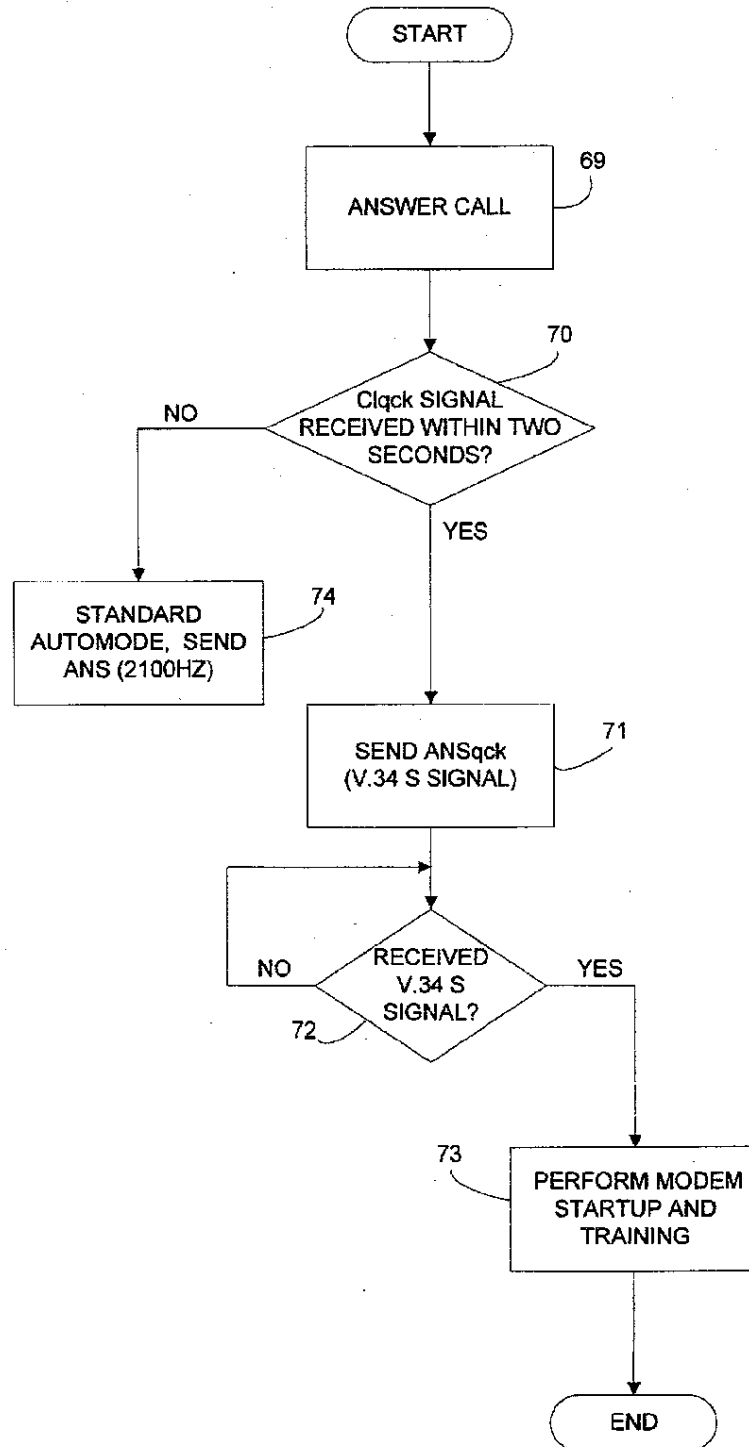
Calling Modem -- Cellular
FIG. 4

U.S. Patent

Dec. 22, 1998

Sheet 4 of 8

5,852,631



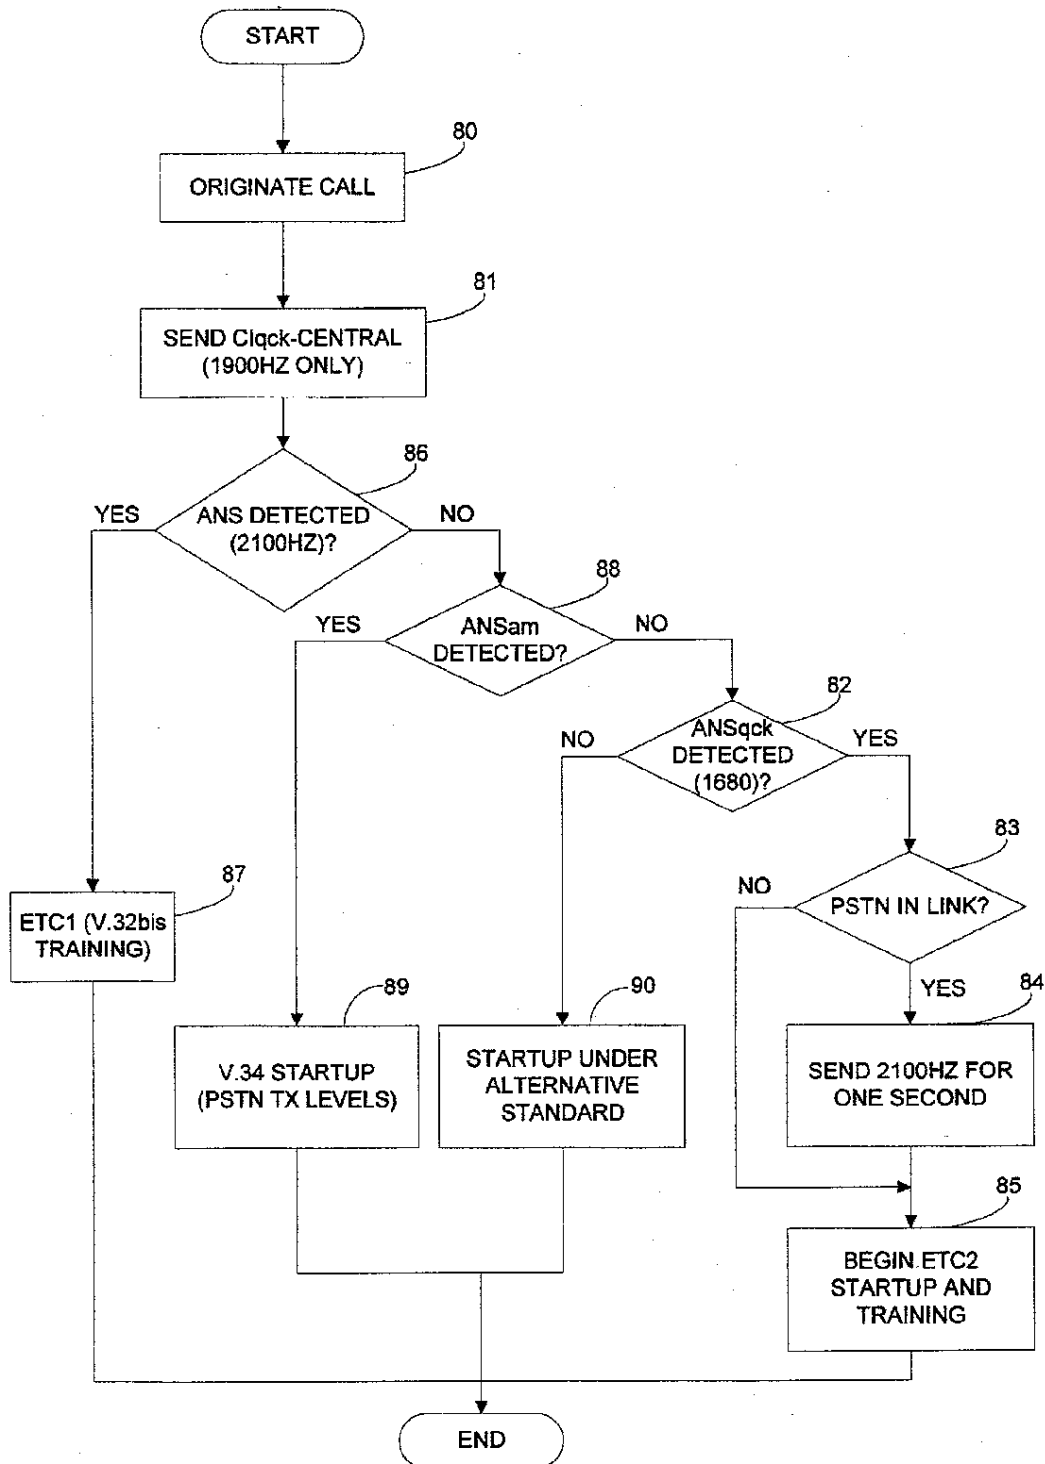
Answer Modem -- Cellular
FIG. 5

U.S. Patent

Dec. 22, 1998

Sheet 5 of 8

5,852,631



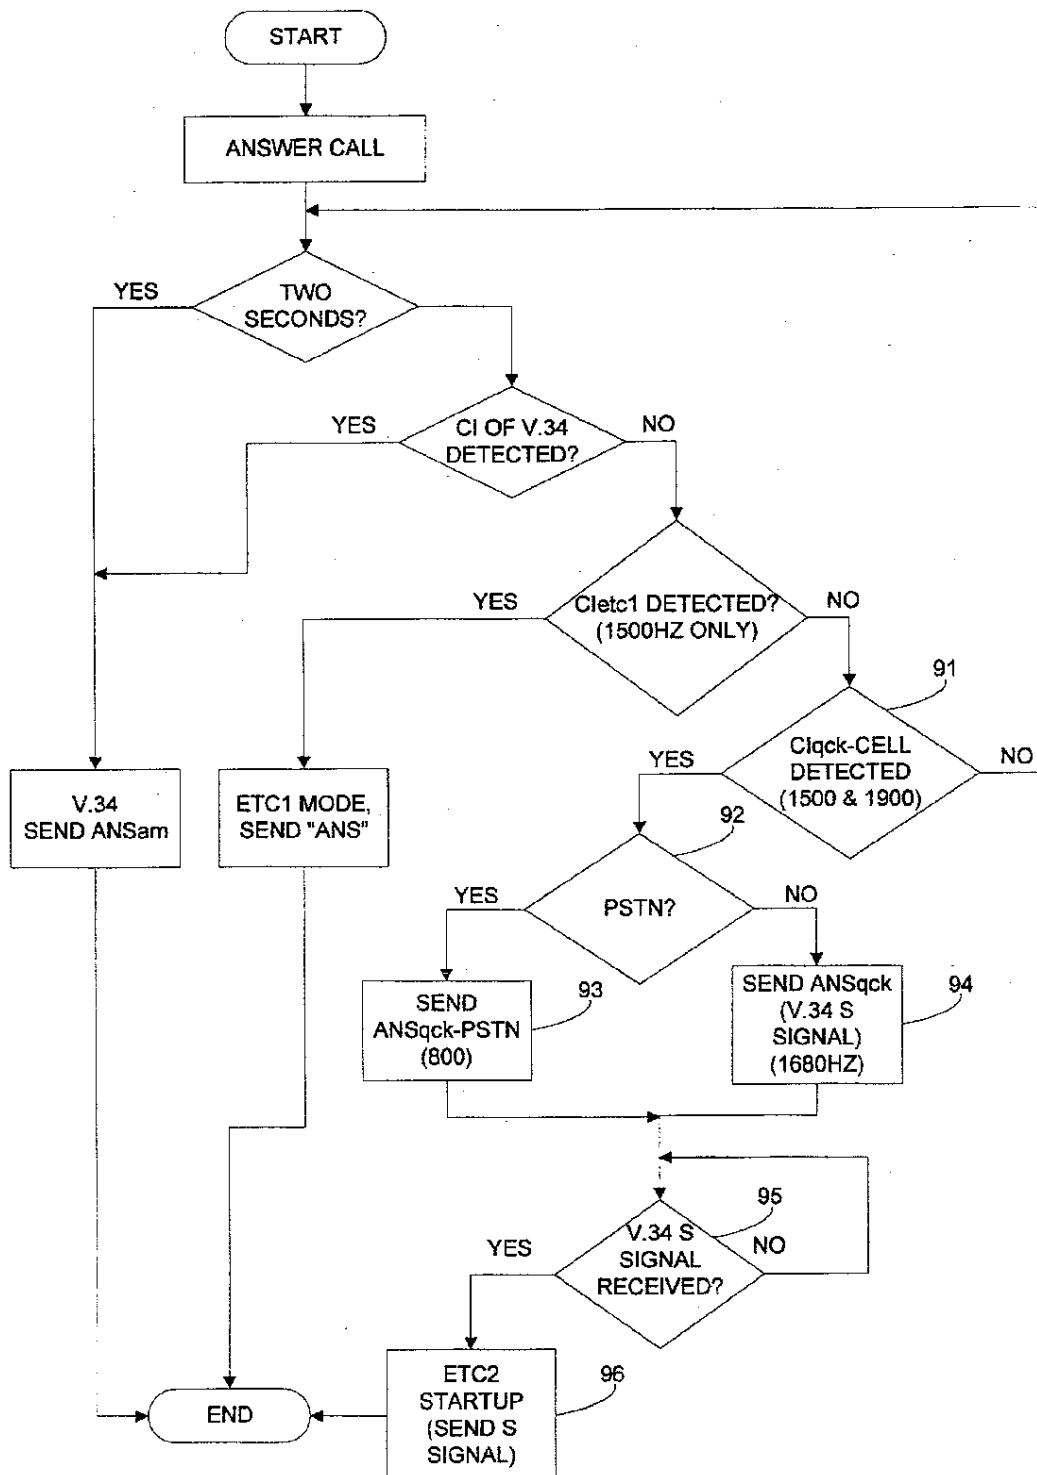
Calling Modem -- Central Site
FIG. 6

U.S. Patent

Dec. 22, 1998

Sheet 6 of 8

5,852,631



Answer Modem -- Central Site
FIG. 7

U.S. Patent

Dec. 22, 1998

Sheet 7 of 8

5,852,631

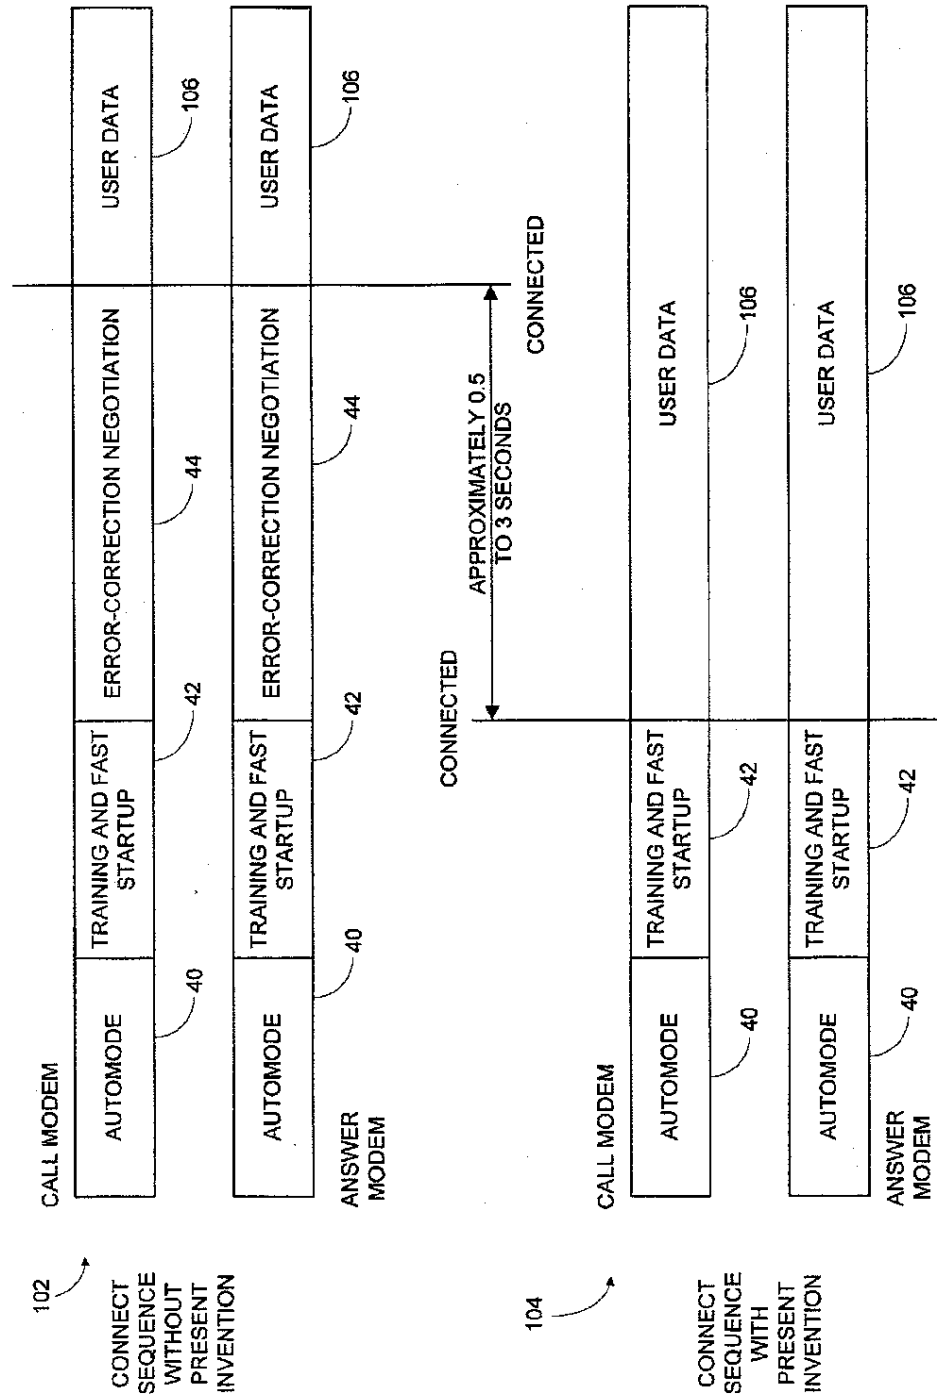


FIG. 8

U.S. Patent

Dec. 22, 1998

Sheet 8 of 8

5,852,631




FIG. 9

5,852,631

1

SYSTEM AND METHOD FOR ESTABLISHING LINK LAYER PARAMETERS BASED ON PHYSICAL LAYER MODULATION

This application claims priority to and the benefit of the filing date of copending and commonly assigned provisional application entitled CELLULAR DATA PROTOCOL FOR QUICK CONNECTION, assigned Ser. No. 60/026,970, and filed Sep. 20, 1996; and copending and commonly assigned provisional application entitled A RAPID START UP PROTOCOL FOR COMMUNICATION BETWEEN A PLURALITY OF MODEMS, assigned Ser. No. 60/022,474, and filed Jun. 21, 1996.

FIELD OF THE INVENTION

The present invention generally relates to data communication protocols, and more particularly, to presetting the link layer parameters per the physical layer modulation in a protocol stack for modems.

BACKGROUND OF THE INVENTION

In an effort to facilitate more reliable and platform independent communication links between remotely located computers, communication protocols are typically organized into individual layers or levels comprising a protocol stack. The lowest layer is designed to establish host-to-host communication between the hardware of different hosts. The highest layer, on the other hand, comprises user application programs which pass customer data back and forth across the communication link. Each layer is configured to use the layer beneath it and to provide services to the layer above it.

Examples of two protocol stacks are the Opened Systems Interconnect (OSI) seven layer model and the Transmission Control Protocol/Internet Protocol (TCP/IP) five layer model. The OSI seven layer model comprises the following layers from lowest to highest: a physical layer, a data link layer, a network layer, a transport layer, a session layer, a presentation layer, and an application layer. When combined, the seven layers form a protocol stack that is designed to provide a heterogeneous computer network architecture. The TCP/IP five layer model comprises the following layers from lowest to highest: a physical layer, a data link layer, a network layer, a transport layer, and an application layer. Of particular relevance to the present invention is the implementation of the physical layer and data link layer in these systems.

The physical layer of the OSI model is the lowest layer and is concerned with establishing the electrical and mechanical connection between two modems. The data link layer is the second lowest layer of the OSI seven layer model and is provided to perform error checking functions as well as retransmitting frames that are not received correctly.

As is well known, a variety of standards exist which govern the protocols for communication between modems. For example, V.21, V.22, V.32, V.32bis, V.34, V.42, and V.42bis, are identifiers of differing communication standards recommended by the International Telecommunications Union (ITU). Each one of these is directed to an aspect of either the physical layer or data link layer of the OSI model.

The ITU Standard V.34 (hereafter referred to as V.34) is intended for use in establishing a physical layer connection between two remotely located computers over the Public Switch Telecommunications Network (PSTN). The V.34 standard includes the following primary characteristics: (1) full and half-duplex modes of operation; (2) echo cancella-

2

tion techniques for channel separation; (3) quadrature amplitude modulation for each channel with synchronous line transmission at selectable symbol rates; (4) synchronous primary channel data signaling rates ranging from 2,400 bits per second to 33,600 bits per second, in 2,400 bit-per-second increments; (5) trellis coding for all data signaling rates; and (6) exchange of rate sequences during start-up to establish the data signaling rate. The features of V.34 are documented in the publicly-available ITU Standard V.34 Specification and are well known by those skilled in the art, and will not be described in detail herein.

Another significant feature of V.34, as it relates to the present invention, is the ability to automode to other V-series modems that are supported by the ITU Standard V.32bis automode procedures. In this regard, V.34 defines signal handshaking that two connecting modems exchange at startup in order to learn the capabilities of the other modem to most efficiently exchange information.

While V.34 achieves efficient and generally high speed communication between two communicating modems, it nevertheless possesses several shortcomings that impede even more efficient operation. One significant shortcoming is the lengthy startup sequence which takes approximately 10-15 seconds. Particularly, for cellular customers, the ability to provide faster connections and faster data rates is particularly desirable since the cellular customer typically pays a charge for each cellular call based primarily on the length of the call and several other factors such as day of the week, time of day, roaming, etc. As a result, new fast connect protocols are being developed that provide for faster and more efficient startup operation based upon the system configuration and the path of the established communication link. An example of one such fast connect protocol is Paradyne Corporation's Enhanced Throughput Cellular 2 Quick Connect™ (ETC2-QC™). In essence, the ETC2-QC™ protocol uses techniques in the physical layer to reduce the physical layer startup time delay to about 1 second.

Of particular relevance to the present invention is the ITU Standard V.42 (hereinafter referred to as V.42). The V.42 standard is intended for use in establishing the error-correcting protocol of the data link layer connection. The V.42 standard includes a detection phase which determines whether both modems are capable of an error-corrected connection, an exchanging identification phase for determining error-correcting parameter values and a link establishment phase for establishing the error-corrected connection. Under normal circumstances, V.42 requires approximately 1-3 seconds to establish an error-corrected connection. While this is relatively small in comparison to the establishment of a physical layer connection under V.34, it can essentially double the connection time when used in conjunction with fast connect modems.

Therefore, a heretofore unaddressed need exists in the industry for a system and method that reduces or eliminates the time required to establish a link layer connection so as to minimize the amount of time for establishing a connection between two modems.

SUMMARY OF THE INVENTION

The present invention overcomes the inadequacies and inefficiencies of the prior art as discussed hereinbefore and well known in the industry. The present invention provides a system and method for establishing a link layer connection between a calling modem having a plurality of possible first physical layer modulations and a plurality of possible link

5,852,631

3

layer connections and an answering modem having a plurality of possible second physical layer modulations and a plurality of possible second link layer connections that comprises the following steps. One step includes establishing a physical layer connection between the calling and the answering modems, wherein the physical layer connection is based on a negotiated physical layer modulation chosen from the first and second physical layer modulations. Another step includes establishing a link layer connection based upon the negotiated physical layer modulation. This link layer connection includes parameters that are preset to default values based upon the negotiated physical layer connection. Thus, the modems are able to avoid the link layer negotiation that essentially all other modems perform, thereby providing a faster and more robust connection.

Other features and advantages of the present invention will become apparent to one with skill in the art upon examination of the following drawings and detailed description. It is intended that all such additional features and advantages be included herein within the scope of the present invention, as defined by the claims.

DESCRIPTION OF THE DRAWINGS

The present invention can be better understood with reference to the following drawings. The elements of the drawings are not necessarily to scale, emphasis instead being placed upon clearly illustrating the principles of the present invention. Furthermore, like referenced numerals designate corresponding parts throughout the several views.

FIG. 1 is a system diagram, illustrating a multi-modem system, wherein a plurality of modems are interconnected among a plurality of communication links;

FIG. 2 is a diagram illustrating the primary handshaking and data exchange sequences between a calling and an answer modem;

FIG. 3 is a timing diagram similar to FIG. 2, illustrating the signal exchange during the automatic mode synchronization sequence of FIG. 2;

FIG. 4 is a software flowchart illustrating the operation of the present invention when the calling modem is a cellular modem;

FIG. 5 is a software flowchart illustrating the operation of the present invention when the answer modem is a cellular modem;

FIG. 6 is a software flowchart illustrating the operation of the present invention when the calling modem is a Central-site modem;

FIG. 7 is a software flowchart illustrating the operation of the present invention when the answer modem is a Central-site modem;

FIG. 8 is a schematic diagram comparing a first connect sequence of two fast connect modems with a conventional link layer connection and a second connect sequence of two fast connect modems with a link layer connection based on the physical layer negotiation in accordance with the present invention; and

FIG. 9 is a block diagram of the modems comprising a data gateway connected to the mobile switching center of FIG. 1.

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENT

The following description is of the best presently contemplated mode of carrying out the present invention. This

4

description is not to be taken in a limiting sense, but is made merely for the purpose of describing the general principles of the invention. Consequently, the scope of the invention should be determined by referencing the appended claims.

The following description is divided into two parts. The first part discloses an example of a fast connect protocol for use in a modem system that is suitable for operating in conjunction with the present invention. It should be noted that the modem system disclosed in the first part is merely illustrative of a system that can benefit from the present invention, as will be evident to those of ordinary skill in the art upon reading the following disclosure. The second part discloses the present invention in the context of the fast connect modem system described in the first part. However, the present invention is equally well suited for application outside the context of the fast connect modem system described herein, for example, with modems that connect slowly.

I. Physical Layer Connection

Turning now to the drawings, FIG. 1 shows a system diagram of a system illustrating multiple modems intercommunicating through a variety of mediums, including cellular and PSTN. Indeed, as previously mentioned, a driving factor in the development of the present invention was to design a system that provided improved reliability in data communication over a data communication link. This goal has been achieved by removing the necessity to perform error-correction negotiation during the connect sequence in modem communication so as to reduce the overall connection time.

As illustrated, a cellular modem system may be disposed for communication with a mobile switching center (MSC) 12. More specifically, a cell 14 includes a portable computer 15 that is connected via a cellular modem 16 to a cellular phone 17, which in turn communicates (wirelessly) with a cell tower 18 that communicates with the MSC 12. It is appreciated that the modem 16 recognizes that it is on the cellular side via a strap or configuration setting, or alternatively by a direct connect sensing of the cellular phone. Therefore, and as will be discussed in more detail below, the modem 16 will know that it is capable of communicating in accordance with a modulation standard capable of performing a fast connect sequence as described below.

The MSC 12 is also connected to a data gateway comprising modems 20 and 24. The modems 20 and 24 are illustrated as connected in a back-to-back configuration and communicating to the MSC 12 over links 22 and 26. As will be appreciated and discussed below, the links 22 and 26 will support different communication protocols, or different modulation standards. The modem pool provides a data gateway for interfacing data calls originating from the modem 16, and thereby, allows cellular specific protocols to be used over the wireless connection between modems 16 and 20, as described below.

By way of definition, a "Central-site" modem is one that is capable of supporting the modulation standard of the present invention, and is not connected to a cellular phone. In this regard, all central-site modems are connected via four-wire connections. Examples which are illustrated in FIG. 1 include a mobile switching center MSC(Cellular) modem 20, an MSC(PSTN) modem 24, an MSC(Single-ended) modem 28, and a PSTN(ETC2-QC™) modem 30—where an MSC modem is one that is connected at the mobile switching center 12. A significant distinction among these various types of modems relates to the startup sequence, which will differ slightly depending upon the type

5,852,631

5

of central-site modem. Preferably, a hardware identifier, such as a DIP switch or a firmware option configurable at modem installation, defines the type of modem for purposes of the startup sequence.

In keeping with the description of FIG. 1, modem 20 is illustrated as an MSC(Cellular) modem that is connected in a back-to-back mode with modem 24, an MSC(PSTN) modem. Modem 20, therefore, is designed to support the ETC2-QC™ modulation protocol and simulate a cellular modem during the initial modem startup routine. Modem 28 is an MSC(Single-ended) modem that, although it may communicate with modems on the PSTN 34, will typically communicate only with cellular modems. Indeed, when communicating with cellular modems, the 2100 Hertz tone, which is typically inserted to disable echo cancellers, is preferably omitted. Advantageously, elimination of this tone achieves a faster and more desirable modem startup.

A PSTN(ETC2-QC™) modem 30 and a standard PSTN modem 32 are connected via PSTN 34 to the MSC 12. The modem 30 is connected to the PSTN 34 via a four-wire connection 35, and modem 32 via a two-wire connection 36. Consistent with the concepts and teachings of the present invention, the four-wire connection 35 facilitates the communication of modem 30 with the cellular modem 16, for example, in the ETC2-QC™ modulation standard. However, as will be appreciated by those of ordinary skill in the art, merely ensuring a four-wire connection 35 alone will not ensure proper system operation in accordance with the present invention. In this regard, such a four-wire connection 35 may nevertheless pass through a two-wire connection, and thus a hybrid converter circuit, at the central office. In this event, echo will be injected into the signal and the abbreviated modulation standard of the present invention may be compromised. There are, however, steps that may be taken to ensure proper operation of the invention. These include, (1) ordering a Direct Inward Dial connection and instructing the phone company to avoid a two-wire connection for that setup; (2) obtaining a direct T1 connection to the Interexchange Carrier (for example, a "1-800" number); and obtaining an ISDN PRI connection, as it will always support four-wire for both call origination and call answer.

By way of illustration, consider a call originated by the computer 15 and cellular modem 16 to the standard PSTN modem 32. The established communication link will pass through the cellular phone 17 to the cell tower 18, through the MSC 12, across link 22 to the MSC(Cellular) modem 20 and to the connected modem 24 via RS-232 connection 38, across link 26 and back through the MSC 12 to the PSTN 34, and ultimately across the two-wire link 36 to modem 32. As will become clear from the description that follows, the cellular modem 16 and the MSC(Cellular) modem 20 will connect and startup in accordance with the fast connect communication protocol described herein. However, since the established communication link that passes from modem 24 to modem 32 passes through a PSTN 34 and a hybrid converter, then the communication protocol of the present invention will not be adequately supported. Accordingly, the modems 24 and 32 will identify this situation and will connect and communicate using an alternative communication protocol supported by both modems and capable of effective transmission across the established link. In this regard, the overall communication link does not realize the fast connection.

Indeed, an aspect of the fast connect protocol described herein is the determination of whether both modems are compatible, in terms of communication protocol, and whether they are connected through a line that passes

6

through a PSTN. If the modems are compatible and the established communication link is outside a PSTN (e.g., cellular to MSC) or is to a PSTN modem with a 4-wire connection that has been configured for supporting a fast connect protocol, then the modems may connect and begin their startup sequence. In this regard, the fast connect communication protocol is designed to be fast as well as robust, and is accomplished by the use of simple tones. The use of such simple tones facilitates the implementation of the automatic mode select to be in the modem's control processor rather than the digital signal processor (DSP) chip.

In addition to the fast connect protocol discussed in below, the fast connect protocol also includes several "fall-back" modulations. More particularly, the modem of the present invention will preferably include Paradyne Corporation's Enhanced Throughput Cellular 1™ (ETC1™), V.34, V.32bis, V.32, and V.22bis modulations. Thus, in the previous example, modems 24 and 32 may communicate using one of these communication protocols. These modulation protocols are documented and will be understood by persons of ordinary skill in the art, and will not be discussed herein. It suffices to say that supporting the above-listed modulation standards greatly enhances the flexibility and versatility of a fast connect modem.

To more particularly describe the initial startup sequence in accordance with the modulation standard of the fast connect modem, reference is made to FIGS. 2 and 3. FIG. 2 illustrates the three principal components of modem exchange or communication. After the cellular modem initiates the call, such that a communication link is established, the modems enter a mode select sequence, referred to herein as automatic mode synchronization 40. During this period, the modems exchange parameters that identify the modems, and thus, their communication protocol. This sequence 40, thus, synchronizes the modems for communication in accordance with the same standard or protocol, such as V.34, V.22, V.22bis, etc.

Once the modems have synchronized their communication protocol, or modulation standard, then they enter a training and startup sequence 42. In a manner known in the art, during this sequence the modems may test the established communication link for noise, bandwidth, etc., in order to determine an appropriate rate for communication. The modems may also operate during this period to train their internal echo cancellers by, for example, ranging the established link of communication. In accordance with a related aspect of the fast connect modems, under certain circumstances the modem training and startup sequence may also be significantly shortened to provide a more robust (both time-shortened and reliable) startup sequence. More particularly, the "circumstances" which provide such a robust startup include communicating modems constructed in accordance with the invention detecting an established link of communication that does not pass through any two-wire connections. The completion of this sequence signifies the establishment of a physical layer connection between two modems.

After the physical layer has been established, the communicating modems enter the information exchange/communication sequence, referred to herein as error-correction negotiation 44, in order to establish the link layer connection. This is of particular relevance to the present invention in that it includes negotiation of a error-correcting protocol such as V.42. During this sequence 44, the modems detect whether they are error-correcting modems and, if so, they negotiate the error-correcting parameters.

Referring now to FIG. 3, the initial automatic mode synchronization 40 is illustrated. As shown, this sequence is

5,852,631

7

executed by exchanging signals between the calling modem and the answer modem. After the calling modem instructs the cellular phone to establish the communication link with the answer modem, it transmits the calling signal Clqck 50. As will be described in more detail in connection with the flowcharts of FIGS. 4-7, this signal may comprise a 1900 hertz tone, or alternatively may comprise a 1500 hertz tone modulated with a 1900 hertz tone. If only a 1900 hertz tone is transmitted as Clqck signal 50, then the answer modem knows that the calling modem is configured as a Central Site, four-wire modem (see FIG. 6). Alternatively, if the Clqck signal includes both 1500 and 1900 hertz components, then the answer modem knows that the calling modem is configured as a cellular modem.

As will be appreciated by those of ordinary skill in the art, other calling signals may be transmitted by the calling modem. For example, calling signals consistent with that of a facsimile transmission, or calling signals consistent with other modem modulation standards, such as V.34, V.32, V.32bis, etc., may be transmitted. Since automatic connection and synchronization to facsimile, and these other modulation standards, are well known it will not be discussed herein. Indeed, the significance of the fast connect protocol is achieved when both the calling modem and the answer modem are capable of communicating in accordance with the fast connect modulation protocol herein described so that through the exchange of tones, the modems are made aware of the possible shortcuts in the fast startup and training sequence 42, and more particularly, in the error-correction negotiation 44.

Once the Clqck signal 50 is received by the answer modem, then the answer modem transmits its response back to the calling modem. The purpose of this answer signal is not only to signal receipt of the calling signal, but also to uniquely identify the answer modem. Again, as is known in the art, this answer signal may comprise ANS or ANSam signals as are known by the V.34 and V.32bis communication protocols. If so, the calling modem will then startup and train 42 and perform error-correction negotiation 44. Significant to the present invention, however, is when the answer signal is ANSqck, which is defined by either a 1680 hertz tone or an 800 hertz tone.

As illustrated in FIG. 4 (assuming the calling modem is a cellular modem), if ANSqck is an 800 hertz tone, then the calling modem knows that the answer modem is configured as a four-wire connection, and can communicate with the calling modem in accordance with the fast connect communication protocol and, in accordance with the present invention, set the error-correction parameters to preset values so as to avoid the necessity of negotiating the parameters. In addition, the 800 hertz ANSqck signals the calling modem that the answer modem is connected to a PSTN 34 (see FIG. 1). Therefore, the calling modem transmits a 2100 hertz tone for approximately one second. This, as is known, serves to pad the initial two second connect period, as required by the FCC for billing purposes. Furthermore, it serves to disable the echo cancellers within the PSTN 34.

If ANSqck is a 1680 hertz tone, which is the center tone of V.34 S signal, then the calling modem knows that the answer modem is configured as a four-wire connection, and can again communicate with the calling modem in accordance with the fast connect communication protocol and, in accordance with the present invention, set the error-correction parameters to preset values so as to avoid the necessity of negotiating the parameters. More significantly, it tells the cellular calling modem that the answer modem is not connected to the PSTN 34. Therefore, both the calling

8

modem and the answer modem can determine that the established communication link is entirely outside the PSTN 34. Accordingly, the Federal Communications Commission (FCC) billing delay need not be inserted. Furthermore, certain assumptions may be made in regard to bandwidth, or transmission quality. For example, the established communication link will not pass through echo cancellers, and as a result, the calling modem need not transmit the 2100 hertz tone. Instead, upon receiving the ANSqck answer signal, the calling modem may immediately enter the modem training and startup sequence 42.

As will be further appreciated by those of ordinary skill in the art, by making certain assumptions regarding the line quality of the established link, the modem training and startup sequence 42 may be shortened. For example, in the preferred embodiment, the system initiates communication by assuming a 9600 baud rate. It has been found that most cellular connections may transmit at this rate, and certain front-end savings may be realized by defaulting to this initial startup rate. Of course, this rate may be increased, or autorated upwardly, in accordance with methods known in the prior art, after the initial startup and training sequence 42 has been completed.

Referring back to FIG. 4, a top-level flowchart is shown, illustrating the automatic mode synchronization of a cellular calling modem constructed in accordance with the fast connect protocol disclosed herein. Once the calling modem has completed transmitting the dialing sequence, it transmits the Clqck signal, which for a cellular calling modem includes modulated 1500 and 1900 hz tones, as indicated in block 60. Once the calling signal has been transmitted, the calling modem will wait to receive the answer signal from the answer modem. In order to exchange data using the modified modulation standard of the present invention, the calling modem looks to receive one of two answer signals. The first valid answer signal as in 1680 hz tone, which is the center tone of the V.34 S signal, as indicated in block 61. This tone signals to the calling modem that the answer modem is not only compatible to transmit in the fast connect modified modulation standard, but further indicates that the answer modem is connected via four wire connections, and does not interconnect to a PSTN. Accordingly, since the calling modem is a cellular modem, then the established communication link does not pass through a PSTN and the initial two second FCC-required delay need not be inserted into the start-up sequence. Moreover, since the entire communication link is four wire, then the modems need not transmit the 2100 hz signal to disable echo cancellers.

A second valid answer signal is an 800 hz tone, as shown by block 62, which also indicates that the answer modem is connected via four wire, and therefore, can communicate in accordance with the fast connect modulation protocol. In addition, the 800 hz tone indicates that the answer modem is connected to a PSTN. Assuming, as previously discussed, that the requisite steps have been taken to ensure that the established communication link does not pass through a two wire connection, then certain savings or efficiencies can be gained during the modem start-up and training sequences (e.g., eliminate echo training since no hybrid circuits are present in the communication link). Nevertheless, the FCC-required delay must be inserted and, therefore, a 2100 hz tone is transmitted at block 63 by the calling modem for a duration of approximately one second. The amount of the 2100 hz tone will "pad" the total modem autotone and startup time to two seconds. This ensures that no customer data is transferred in the first two seconds (which meets FCC requirements). Thereafter, calling modem proceeds with the modem training and start-up sequence at block 64.

5,852,631

9

If neither of the foregoing answer signals are received, then the system operates to determine whether another valid answer signal has been transmitted from the answer modem. The step of block 65 broadly designates this function. It should be appreciated that well known answer signals such as ANS or ANSam may be transmitted by the answer modem and, if received, the calling modem may synchronize to the appropriate modulation standard, as indicated in block 66. Although not separately designated in the figure, it should be further appreciated that if no valid answer signal is received by the calling modem within a given period of time, the calling modem will time out and abort the attempted communication. Also, and as illustrated at block 67, the calling modem will abort the attempted communication if a busy signal is received.

FIG. 5 shows a top-level flowchart illustrating the operation of a cellular answer modem constructed in accordance with the fast connect communication protocol described herein. Once the communication link has been established and the call answered at block 69, the answer modem looks to detect the Clqck calling signal, as indicated by block 70. In the presently described fast connect protocol, cellular to cellular modem communications are not supported. Therefore, a cellular answer modem will assume that a calling modem transmitted a Clqck signal will transmit only a 1900 hz tone rather than the modulated 1500 and 1900 hz tones. Having said this, it should be appreciated that cellular-to-cellular communications could be supported.

In keeping with the description of FIG. 5, once the answer modem has received the Clqck calling signal, it transmits the ANSqck answer signal at block 71. It then waits for the calling modem to enter the modem start-up and training sequence. This sequence is identified by receiving the S signal as assigned by the V.34 modulation standard, as indicated by the decision block 72. Once this signal is received, then the answer modem will transmit back to the calling modem the appropriate S signal, so as to initiate the startup and training sequence 42.

Alternatively, if the answer modem, within a period of two seconds, has not received Clqck calling signal, then it will proceed with the start-up sequence in accordance with an alternative modulation standard. This, therefore, assumes that the modified communication protocol of the present invention is not supported by the calling modem, and the answer modem will typically respond to the calling signal of an alternative communication signal by transmitting a 2100 hz tone, as indicated in block 74.

Referring now to FIG. 6, a software flowchart illustrating the top-level operation of a central site calling modem is shown. As depicted, the calling modem originates the call and establishes a communication link at block 80. Once the communication link is established, the calling modem transmits the Clqck calling signal at block 81, which in the case of a central site calling modem comprises a 1900 hz signal tone. If the 1680 hz ANSqck answer signal is detected at block 82, then the calling modem recognizes the answer modem as one capable of transmitting pursuant to the fast connect communication protocol. Thereafter, the calling modem must determine the network configuration of the established communication link, as indicated in block 83. That is, the central site calling modem will determine whether the established communication link passes through a PSTN or not. If it is determined that the established link passes through a PSTN, then, as in the case of the cellular calling modem, the calling modem transmits a 2100 hz signal for approximately one second at block 84. Thereafter, the calling modem enters the modem start-up and training sequence, as indicated in block 85.

10

Alternatively, if the calling modem detects the ANS answer signal (2100 hz) at block 86, then it communicates with the answer modem using the ETC1™ communication protocol and the V.32bis training, as indicated by block 87. If the ANSam answer signal is detected at block 88, then the modem will startup in standard V.34 mode at block 89, which is well known in the art and therefore not described herein. The modem will also monitor for ANSqck at block 82, which in this example is a 1680 Hz tone. If this is not detected, then the modem will startup under an alternate low speed standard at block 90, which is well known in the art and therefore not described herein. If ANSqck is detected, then the modem will operate differently depending on whether it is connected to the PSTN network or not, as indicated by block 83. The modem will know whether it is connected to the PSTN via a configuration option which was set at installation. If connected to the PSTN, then the modem will transmit a 2100 Hz tone for one second at block 84 then proceed to the ETC2™ training sequence at block 85. If the modem is not connected to the PSTN at block 83, then it can proceed directly to the ETC2™ training sequence at block 85, avoiding the additional one second of startup shown at block 84.

Reference is now made to FIG. 7, which is a software flowchart illustrating the top-level operation of a central site answer modem. As illustrated in the flowchart, and in accordance with the presently disclosed fast connect protocol, when the answer-site modem is a central-site modem, it assumes that any transmissions made in accordance with the modulation standard with the present invention will be via a communication link with a cellular calling modem. Therefore, block 91 indicates detection the Clqck calling signal in the form of a modulated 1500 and 1900 hz tones, as transmitted by cellular calling modem. If the Clqck calling signal is detected, then the answer modem determines the network configuration at block 92. More specifically, the answer modem determines whether the established communication link passes through a PSTN or not. In the event that the established link does in fact pass through a PSTN, then the answer modem will transmit an 800 hz ANSqck answer signal at block 93. As illustrated in FIG. 4, this instructs the calling modem to transmit the 2100 hz tone. Alternatively, the answer modem will transmit the 1680 hz tone, which instructs the calling modem to proceed directly with the modem start-up and training sequence at block 94. Thereafter, the answer modem will await transmission of the S signal in accordance with the V.34 start-up sequence, as indicated by block 95. Thereafter, the answer modem will respond by transmitting the S of the V.34 start-up, as indicated by block 96. Since the V.34 start-up sequence is well-known in the art, it would not be described herein.

The remainder of the flowchart depicted in FIG. 7 illustrates the central-answer modem operation and connects sequence in accordance with alternative standards that are well-known in the prior art and need not be discussed herein.

Accordingly, at the completion of the automatic mode synchronization sequence 40 (FIG. 2), the modems enter into a training and start-up sequence 42. As mentioned above, in the training and startup sequence 42 the modems test the established communication link for noise, bandwidth, etc., in order to determine the appropriate rate for communication. This is performed using the modulation scheme determined in the automatic mode synchronization sequence 40 as illustrated in FIGS. 4, 5, 6, and 7. For purposes of the following discussion, it is assumed that the call and the answer modems are capable of communicating

5,852,631

11

with one another using a fast connect protocol. Consequently, since the call modem knows what type of modem it is and what type of modem the answer modem is, certain shortcuts can be taken during the training and startup sequence 42 so as to reduce the overall connection time. Specifically, the modems can default to preset values that eliminate the need for probing, ranging and half-duplex training. Thus, the modems merely perform a special full-duplex training mode during the training and start-up sequence 42 which results in a much faster connection.

Particularly, the probing and ranging sequences are bypassed and the file parameters are assumed in ITU Standard V.8, INFO0, and INFO1. As an example, in ITU Standard V.8, the data call, the LAPM and the full-duplex training parameters are preset to default values if the tones exchanged during automode sequence indicate that both modems are capable of fast connect operation. Further, in the INFO sequences, the 4 point train, 2800 L symbol rate, the power level drop, and preemphasis filter can also be preset to default values. Thus, the ITU Standard V.8 and INFO sequences are eliminated.

At the completion of the training and start-up sequence 42, the modems have established a physical layer connection and are ready to establish the second layer connection, referred to as the link layer connection, via an error-correction negotiation sequence 44 in accordance with the present invention, as disclosed below.

II. Link Layer Connection

The link layer is the second layer of the ISO model protocol stack and includes negotiating and establishing an error-correcting connection such as with ITU Standard V.42 or Microcom Networking Protocol (MNP). The link layer connection follows the physical layer connection and uses the physical layer in establishing the error-corrected connection. It is noted, however, that conventional wisdom to date has maintained the link layer connection be independent of the physical layer connection when establishing a connection between two modems. In contrast, the present invention establishes the link layer connection based upon the modulation chosen in the physical layer connection during the automatic mode synchronization sequence 40 (FIG. 2). Thus, the steps for establishing an error-correcting protocol are eliminated and the link layer connection is established substantially instantaneously upon the completion of the physical layer negotiation. This not only reduces the amount of time required to establish a connection between two modems, it makes the connection more robust by removing the necessity of performing additional handshaking that, if corrupted for whatever reason, will result in a disconnect or call connect failure.

By way of example, the ITU Standard V.42 (hereafter referred to as V.42) comprises a detection phase, and exchange identification (XID) phase, and a link establishment phase, all of which are briefly discussed below. A more detailed explanation of V.42 can be found in the publicly-available ITU (CCITT) Recommended Standard V.42 documentation.

The detection phase is provided to determine whether the answer modem supports an error-correcting protocol. This phase is designed to avoid the potential disruptions to the answer DTE that could occur if the calling modem immediately enters the XID phase and the answering modem was not capable of an error-correcting communication. However, the detection phase is optional and may be disabled. If the call modem determines that the answering modem does not support a V.42 error-correcting protocol, there are often

12

times fall-back error protocols provided by the calling modem, such as in the case of V.42, where MNP is provided as a fall-back error-correcting protocol. Alternatively, if the answer modem does not support V.42 nor MNP, then no error-correcting protocol is established and a connect message is issued by the modems to their respective digital terminal equipment so that user data can be transmitted between the two modems.

The XID phase is provided for the negotiation of the error-correcting parameter values. These parameters essentially govern the error-correcting operation of the modems once the connection is established. As with the detection phase, the XID phase may be omitted if default parameter values are acceptable. For example, the following are provided as the default parameters values in the V.42 standard: Standard Reject, 16 bit FCS (Frame Check Sequence), V.42bis compression disabled, Frame Length (N401) of 128 octets, and Window Size (k) of 15 frames. However, the default settings are more often than not undesirable because, for example, most modems wish to negotiate Selective Reject, V.42bis data compression, and longer Frame Lengths and Window Sizes.

Lastly, the link establishment phase is provided for actually making the error-corrected connection between the two modems. In V.42, this is implemented via a set asynchronous balanced mode extended (SABME) command. The SABME command is used to place the addressed error-corrected entity (i.e., the answering modem) into the connected state. The error-correcting entity then confirms acceptance of the SABME command by the transmission of an unnumbered acknowledgment (UA) response. By acceptance of this command, the error-corrected connection is essentially established and the modems then send a connect message to their respective data terminal equipment, such as computer 15 (FIG. 1).

Unlike the detection phase and the XID phase, the link establishment phase is not optional and must be performed under V.42. Thus, in a best case scenario, only the link establishment phase is performed, which takes approximately 0.5 seconds. If all three phases are performed, then the link layer connection may take three or more seconds.

Therefore, establishing an error-corrected connection with V.42 can take up to three seconds, depending on what defaults are set in the system. While this amount of time does not seem significant relative to the time required for establishing a physical layer connection via V.34 modulation (e.g., approximately 10-15 seconds), it is considerably more noticeable when a fast connect protocol is utilized that can establish a physical layer connection in about 1 second. Thus, when using a fast connect protocol, the error-correction negotiation can easily double the connect time, not to mention introduce a greater opportunity for failure by requiring additional handshaking.

Accordingly, the present invention enables an error-corrected connection without having to perform the steps described above with regard to V.42, or those steps associated with other error-correcting protocols as known in the art. The present invention achieves this by presetting the XID phase parameters to default values that are based upon the negotiated physical layer connection. Therefore, when two multi-mode modems negotiate a physical layer connection, the link layer connection can be immediately established based upon the negotiated physical layer modulation. For example, in the embodiment described above in Section I, the exchange of tones in the mode synchronization sequence 40 indicates to each modem the type of modem it

5,852,631

13

is communicating with, and therefore, certain assumptions can then be made regarding the error-correction negotiation sequence 44 so as to eliminate the steps normally performed to establish an error-corrected connection. In the preferred embodiment of the present invention with the V.42 standard, the following parameters are set to the indicated default values when the two modems are capable of the fast connect sequence described above: Selective Reject, 16 bit FCS, 64 bit Maximum Frame Size (transmit and receive directions), 8 Frame Window Size (transmit and receive directions), V.42bis enabled, and 1,024 bit dictionary (transmit and receive directions). It should be noted, however, that one of ordinary skill in the art would recognize that these default values are merely illustrative settings and that different default values can be used. Moreover, each different type of connect sequence would preferably have its own set of default values. If it is determined by the modems in the mode synchronization sequence 40 that one or the other is not capable of a fast connect as described above, then the modems essentially fallback and perform an alternative error-correction sequence such as the recommended ITU Standard V.42 error-correction sequence.

With reference to FIG. 8, a graphical illustration is provided of two fast connect modems in a first connect sequence 102 where error-correction negotiation is performed without the present invention and a second connect sequence 104 where the error-correction negotiation is performed with the present invention. As shown, following the mode synchronization sequence (also referred to as autumode) 40 and the training and start-up sequence 42, the connect sequence 102 performs error-correction negotiation 44 which essentially doubles the time required for a connection to be established so as to allow user data 106 to be exchanged. In comparison, the connect sequence 104 in accordance with the present invention is able to establish a connection in essentially half the time by eliminating the error-correction negotiation 44. Thus, by establishing the error-correction parameters to default values in accordance with the type of physical error-connection determined by the autumode sequence 40, a faster and more reliable connection is established.

Regarding the implementation of the present invention, FIG. 9 generally illustrates the components of MSC (cellular) modem 20 and MSC(PSTN) modem 24 which implement the data gateway. The MSC(cellular) modem 20 comprises a digital signal processor (DSP) 112, a control processor 114, and a DTE interface 116. Likewise, the MSC(PSTN) modem 24 comprises a DSP 118, a control processor 120, and a DTE interface 122. The DTE interface 116 of the MSC(cellular) modem 20 interfaces with the DTE interface 122 of the MSC(PSTN) modem 24 via the connection 38, which can be implemented by any suitable interconnecting device such as, but not limited to, an Electronic Industry Association (EIA) standard RS-232 crossover or a backplane bus between the modems. As shown in FIG. 9, each modem 20, 24 is configured essentially the same, and thus, they operate in essentially the same manner. However, each modem is provided with operating code which is stored in a memory device 124, 124' provided with the central processor 114, 120, respectively, though additional memory can also be provided and connected to the control processor 114, 120, if necessary. In the context of the present disclosure, a memory device is a computer readable medium that is embodied in an electronic, magnetic, optical or other physical device or means that can contain or store a computer program, such as the operating code for the modem 20, 24, for use by or in connection with a computer

14

related system or method. The operating code includes control logic that controls, among other things, the type of modulation and error correction techniques utilized which is dependent upon whether the modem is used for cellular or land-line connections. Accordingly, the control processor 114, 120 operates on, or executes, the operating code that is in memory device 124, 124' and configured for implementing the present invention so as to control the operation of modem 20, 24.

The foregoing description has been presented for purposes of illustration and description. It is not intended to be exhaustive or to limit the invention to the precise forms disclosed. Obvious modifications or variations are possible in light of the above teachings. The embodiment or embodiments discussed were chosen and described to provide the best illustration of the principles of the invention and its practical application to thereby enable one of ordinary skill in the art to utilize the invention in various embodiments and with various modifications as are suited to the particular use contemplated. All such modifications and variations are within the scope of the invention as determined by the appended claims when interpreted in accordance with the breadth to which they are fairly and legally entitled.

Wherefore, the following is claimed:

1. A method for establishing a link layer connection between a calling modem having a plurality of possible first physical layer modulations and a plurality of possible link layer connections and an answering modem having a plurality of possible second physical layer modulations and a plurality of possible second link layer connections, comprising the steps of:

establishing a physical layer connection between said calling and said answering modems, wherein said physical layer connection is based on a negotiated physical layer modulation chosen from said first and second physical layer modulations; and

establishing said link layer connection based upon said negotiated physical layer modulation.

2. The method of claim 1, wherein said negotiated physical layer modulation is a fast connect modem modulation.

3. The method of claim 1, wherein said link layer connection is an error-correcting protocol.

4. The method of claim 1, further comprising the step of presetting link layer parameters of said link layer connection to default settings based on said negotiated physical layer modulation.

5. The method of claim 3, wherein said error-correcting protocol includes parameters that are set to pre-defined settings based on said negotiated physical layer modulation.

6. A system for establishing a link layer connection between a calling modem having a plurality of possible first physical layer modulations and a plurality of possible link layer connections and an answering modem having a plurality of possible second physical layer modulations and a plurality of possible second link layer connections, comprising:

means for establishing a physical layer connection between said calling and said answering modems, wherein said physical layer connection is based on a negotiated physical layer modulation chosen from said first and second physical layer modulations; and

means for establishing said link layer connection based upon said negotiated physical layer modulation.

7. The system of claim 6, wherein said negotiated physical layer modulation is a fast connect modem modulation.

8. The system of claim 6, wherein said link layer connection is an error-correcting protocol.

5,852,631

15

9. The system of claim 6, further comprising means for presetting link layer parameters of said link layer connection to pre-defined settings based on said negotiated physical layer modulation.

10. A computer program product having a computer 5
readable medium including computer program logic recorded thereon for use in a calling modem for establishing a link layer convention between said calling modem having a plurality of possible first physical layer modulations and a plurality of possible link layer connections and an answering 10
modem having a plurality of possible second physical layer

16

modulations and a plurality of possible second link layer connections, comprising:

logic for establishing a physical layer connection between said calling and said answering modems, wherein said physical layer connection is based on a negotiated physical layer modulation chosen from said first and second physical layer modulations; and

logic for establishing link layer connection based upon said negotiated physical layer modulation.

* * * * *

UNITED STATES PATENT AND TRADEMARK OFFICE
CERTIFICATE OF CORRECTION

PATENT NO. : 5,852,631
DATED : December 22, 1998
INVENTOR(S) : Robert Earl Scott

It is certified that error appears in the above-identified patent and that said Letters Patent is hereby corrected as shown below:

Column 14, line 54, delete the first instance of "a" and replace with -- an --;

Column 15, line 8, delete "convention" and replace with -- connection --; and

Column 16, line 8, after "establishing" insert -- said --.

Signed and Sealed this
Eighteenth Day of May, 1999

Attest:



Q. TODD DICKINSON

Attesting Officer

Acting Commissioner of Patents and Trademarks



US005719858A

United States Patent [19]
Moore

[11] **Patent Number:** 5,719,858
[45] **Date of Patent:** Feb. 17, 1998

[54] **TIME-DIVISION MULTIPLE-ACCESS METHOD FOR PACKET TRANSMISSION ON SHARED SYNCHRONOUS SERIAL BUSES**

[75] **Inventor:** Wayne T. Moore, Clearwater, Fla.

[73] **Assignee:** Paradyne Corporation, Largo, Fla.

[21] **Appl. No.:** 509,309

[22] **Filed:** Jul. 31, 1995

[51] **Int. Cl.⁶** H04B 7/212

[52] **U.S. Cl.** 370/347; 370/349; 370/442; 370/458; 370/468; 340/825.5

[58] **Field of Search** 370/349, 331, 370/335, 355, 442, 462, 461, 460, 463, 352, 353, 468, 447, 405, 458, 347; 340/825.5

[56] **References Cited**

U.S. PATENT DOCUMENTS

4,445,213	4/1984	Baugh et al.	370/405
4,611,321	9/1986	Gabrielli et al.	370/352
4,750,168	6/1988	Trevitt	370/462
4,763,321	8/1988	Calvignac et al.	370/352
5,124,983	6/1992	Landez et al.	370/447
5,463,624	10/1995	Hogg et al.	370/461

Primary Examiner—Wellington Chin


Assistant Examiner—Melissa Kay Carman

Attorney, Agent, or Firm—Thomas, Kayden, Horstemeier & Risley

[57] **ABSTRACT**

A plurality of packet data sources, e.g., packet application modules, and synchronous data sources, e.g., synchronous application modules, are coupled to the same TDM bus for communicating data to a network access module. A portion of the bandwidth, or time-slots, of the TDM bus is allocated as a multiple-access packet channel that is shared among the packet application modules. As a result, the network access module receives a single, continuously multiplexed, packet stream for transmission to an opposite endpoint. Packet application modules on the TDM bus contend for this multiple access packet channel when transmitting to the opposite endpoint. In the receiving direction, each packet application module accepts the entire received packet stream from the network access module and may either filter the packets using their address field or may transparently forward the data to a packet service.

32 Claims, 4 Drawing Sheets



U.S. Patent

Feb. 17, 1998

Sheet 1 of 4

5,719,858

FIG. 1
PRIOR ART




FIG. 2
PRIOR ART


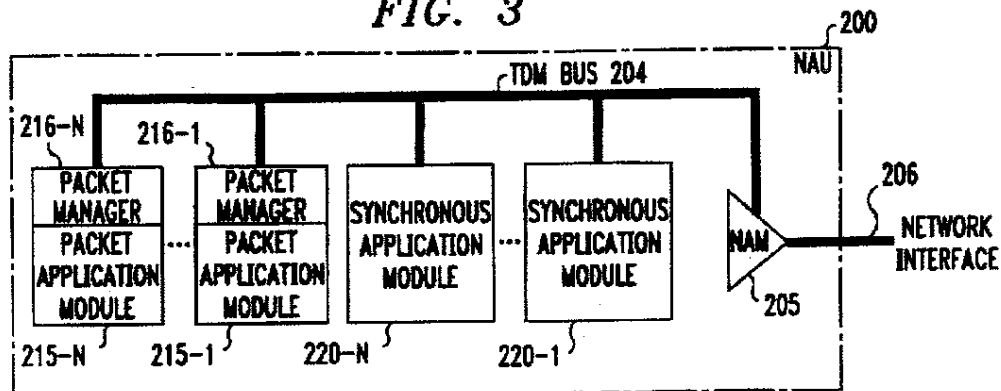


FIG. 3



U.S. Patent

Feb. 17, 1998

Sheet 2 of 4

5,719,858

FIG. 4





FIG. 5



U.S. Patent

Feb. 17, 1998

Sheet 3 of 4

5,719,858

FIG. 6





FIG. 7



U.S. Patent

Feb. 17, 1998

Sheet 4 of 4

5,719,858

FIG. 8A

INITIALIZE




FIG. 8B

REQUEST TO TRANSMIT A PACKET





FIG. 8C

WAIT FOR PHOLD



5,719,858

1

TIME-DIVISION MULTIPLE-ACCESS METHOD FOR PACKET TRANSMISSION ON SHARED SYNCHRONOUS SERIAL BUSES

BACKGROUND OF THE INVENTION

The present invention relates to data communications, and more particularly, to communications systems that have channelized network access, and may transport both synchronous data and variable-bit-rate data such as frame relay (hereafter referred to as packet data), in a time-division multiplexed format.

Communications equipment horn as a "network access unit" (NAU) typically provides frame-relay-type services between a local communications network and a network facility, like a T1 facility. The NAU manages the flow of data between the local communications network and the network facility in both directions. To provide the most flexibility, it is preferable that the NAU support two types of data: synchronous data and packet data. For example, the support of synchronous data provides the ability to make telephone, i.e., voice, calls, while the support of packet data provides the ability to interwork with public network packet services. However, the asynchronous nature of packet data at the logical level combined with the requirements of synchronous data causes design tradeoffs in both the complexity and cost of an NAU.

One prior art approach of designing an NAU to support both synchronous data and packet data is shown in FIG. 1. NAU 100 includes network access module 115-1 105, synchronous application modules 120-1 to 120-n, packet application modules 115-1 to 115-n, and packet manager 110. NAM 105 provides the interface between time-division-multiplexing (TDM) bus 104 and network facility 106, which is representative of a T1 facility. The synchronous application modules couple synchronous data equipment (not shown), e.g., telephone equipment, to NAM 105 via TDM bus 104. The packet application modules couple packet data equipment (not shown), e.g., a data terminal, to packet manager 110. The latter is a common resource module that performs internal aggregation in one direction, and distribution in the other direction, of packet streams between NAM 105 and each packet application module, via TDM bus 104 and wideband packet buses 116-1 to 116-n, respectively. In this context, each wideband packet bus is transmitting packet data in parallel, e.g., a "byte" at a time, in contrast to TDM bus 104, which is a serial bus, i.e., it transmits data a bit at a time. As a result of coupling the packet data to the TDM bus, packet manager 110 provides a single multiplexed packet stream to NAM 105 for transmission across the network interface. It should be understood that the network interface bandwidth is "channelized" and, in the context of packet data, expects a single multiplexed packet stream.

NAU 100 is representative of a mixed TDM and packet NAU architecture. In this approach, NAU 100 provides a TDM bus in conjunction with one or more packet buses which taken together provide more bandwidth than is required to support the network interface. This additional bandwidth is used to support multiple point-to-point packet connections. Packet manager 110 not only aggregates the packet data, as mentioned above, but also allocates a fixed amount of the TDM bandwidth to the packet application modules.

Unfortunately, the instantaneous, or peak, data rate of all outbound packet streams taken together may be greater than the "fixed amount of TDM bandwidth" allocated for packet

2

data on the network interface. These peak data rates create a large demand on both the overall packet bus capacity and on the packet handling requirements of packet manager 110. For example, once the packet application modules exceed their allocated network interface bandwidth, packet manager 110 must take steps to prevent the loss of any packet data. These steps include buffering the packet data, which may require a buffer of considerable size to support all of the packet application modules, and, perhaps, flow control to throttle the packet traffic. As a result, the complexity of packet manager 110 increases not only with the number of packet application modules that packet manager 110 must support but also with the respective data rate requirements of these packet application modules.

Another prior art approach is illustrated in FIG. 2, which is similar to FIG. 1 except that separate point-to-point wideband packet buses have been replaced with separate TDM channels between each packet application module and the packet manager. In particular, packet application modules 165-1 to 165-n are coupled to TDM bus 154, along with packet manager 170. Each packet application module communicates data to, and from, packet manager 170 in a separate TDM channel as represented by lines 171 and 172. Like the description above for FIG. 1, packet manager 170 aggregates the packet traffic to NAM 155 via a TDM channel, as represented by line 173, to create a single multiplexed packet stream.

Unfortunately, this approach has the above-mentioned problems with respect to the packet manager and, in addition, drives up the bandwidth requirement of the TDM bus. For example, the TDM bus must now support the peak rate of each packet application module on each TDM channel in addition to a TDM channel for the aggregate packet data stream provided by packet manager 170. Conversely, if the TDM channels are limited to a fixed fraction of the allocated network interface bandwidth, a single packet application module would never be able to peak near the full rate when other packet application modules have little or no traffic.

SUMMARY OF THE INVENTION

I have realized an alternative approach to the design of TDM-based equipment that supports both synchronous data and packet data and, in addition, provides an efficient substrate for packet handling. In particular, multiple packet data sources share a single TDM channel. As a result, no central packet manager is required to aggregate the packet data.

In an embodiment of the invention, a plurality of packet data sources, e.g., packet application modules, and synchronous data sources, e.g., synchronous application modules, are coupled to the same TDM bus for communicating data to a network access module. In particular, a portion of the TDM bandwidth allocated to packet data is treated as a "multiple-access packet channel." This allows packet application modules on the TDM bus to share, and contend for, the entire TDM bandwidth allocated to packet data. Each packet application module includes its own TDM bus interface and the network access module receives a single, continuously multiplexed, packet stream for transmission to an opposite endpoint. In the receiving direction, each packet application module accepts the entire received packet stream from the network access module and either filters the packets using their address field or transparently forwards the packet data to a packet service.

In a feature of the invention, a contention scheme for accessing the "multiple-access packet channel" is described

5,719,858

3

that avoids packet collisions, maximizes bandwidth efficiency, and provides for interframe High-level Data Link Control (HDLC) flags.

This invention provides the following advantages: no central packet manager is required to synchronize packet data to the TDM bus; packet sources share all of the TDM bandwidth allocated to packet data resulting in maximum efficiency; there is no additional overhead required for packet addressing on the bus; packet buffering is distributed across the bus, rather than being fixed in a central location; and the system has "modularity" and can quickly grow simply by adding additional packet application modules.

BRIEF DESCRIPTION OF THE DRAWING

FIG. 1 shows a block diagram of a prior art mixed TDM and packet network access unit;

FIG. 2 shows a block diagram of another prior art time-division-multiplexing-only network access unit;

FIG. 3 shows an illustrative block diagram of a time-division-multiplexing-only network access unit in accordance with the principles of the invention;

FIG. 4 shows a block diagram of a packet application module in accordance with the principles of the invention;

FIG. 5 shows a representation of a sequence of time-division-multiplexing frames in accordance with the principles of the invention;

FIG. 6 is an illustration of a packet arbitration scheme for use in the network access unit of FIG. 2 in accordance with the principles of the invention;

FIG. 7 is an illustration of counting bit time-slots in the "multiple-access packet channel," and

FIGS. 8A-8C are illustrative flow diagram of a method in accordance with the principles of the invention for use in the packet application module of FIG. 4.

DETAILED DESCRIPTION

An illustrative block diagram of an NAU embodying the principles of this invention is shown in FIG. 3. NAU 200 provides access to a T1 facility for frame relay services as well as synchronous data transport. For simplicity, data endpoints, synchronous or packet, are not shown. NAU 200 includes network access module (NAM) 205, synchronous application modules 220-1 to 220-n, and packet application modules 215-1 to 215-n. NAM 205 provides the interface between TDM bus 204 and network facility 206, which is representative of a T1 facility. The synchronous application modules couple synchronous data equipment (not shown), e.g., telephone equipment, to NAM 205 via TDM bus 204, as is known in the art. In accordance with the inventive concept, multiple packet application modules now share a single TDM channel (described below). Each of the plurality of packet application modules couple packet data equipment (not shown), e.g., a data terminal, to TDM bus 204, and the packet manager is eliminated. Indeed, the function of the packet manager is now distributed among the various packet application modules that created the need for it in the first place. This distribution of the packet manager is represented by the inclusion of packet managers 216-1 through 216-n in the respective packet application modules.

A block diagram of illustrative packet application module 215-n is shown in FIG. 4. Other than the inventive concept, the components of packet application module 215-n are well-known and will not be described in detail. Packet application module 215-n includes packet interface processor 305, buffer 315, and packet/TDM interface 310. The

4

latter is an "application specific integrated circuit" (ASIC), which is a programmable large-scale integrated circuit device that is dedicated to performing a specific function, or application, which in this case is described below. Packet interface processor 305 communicates packet data between packet/TDM interface 310 and line 304. The latter is representative of any one of a number of facilities for coupling packet application module 215-n to a packet system. For example, line 304 could be a local area network, or a dedicated facility to a "router" or a packet data terminal. Packet interface processor 305 performs packet handling, e.g., it provides the physical and link layer connections for packet transmission as known in the art, e.g., it checks for addresses, errors, etc., on the packets. Packet data transmitted to a far-end packet endpoint is provided from packet interface processor 305 to packet/TDM interface 310 via line 307, buffer 315, and line 309. It is assumed that lines 307 and 309 are representative of wideband data buses as known in the art. In the other direction, packet data from a far-end packet endpoint is received by packet interface processor 305 from packet/TDM interface 310 via line 306. Packet interface processor 305 receives a BUFFER FULL signal via line 314 from buffer 315, which provides a BUFFER NOT EMPTY signal via line 321 to packet/TDM interface 310. In accordance with the principles of the invention, packet/TDM interface 310 synchronizes packet data retrieved from buffer 315 for insertion into an appropriate time slot on TDM bus 204. The latter is shown as actually comprising two TDM buses. TDM bus 204-o is used for "outbound" traffic through NAM 205 to network interface 206. Conversely, TDM bus 204-i is used for "inbound" traffic from the network. Typically, in the art, TDM bus 204-i and 204-o are symmetrical, e.g., if time-slot 0 is used for status and control information on the "inbound" TDM bus, time-slot 0 of the "outbound" TDM is similarly used for status and control information.

The storage size of buffer 315 is determined empirically and is a function of the type of packet equipment and its data rate. For example, if a packet application module is interfacing to a router, as known in the art, a router may communicate packet data on the order of 128 Kb/sec. Consequently if the packet application module cannot, for the moment, transmit a packet from the router to the TDM bus, the packet application module can perform flow control with the router with a concomitant minimal amount of buffering required on the packet application module because of the low data rate from the router. Conversely, if a packet application module is interfacing to a LAN, which typically communicates data at a higher speed, e.g., 10 Mb/sec., a larger amount of buffering may be required on the packet application module. Thus, depending on the packet equipment, each packet application module may have different buffering requirements—but each packet application module has less buffer requirements than the packet manager module of the prior art.

In accordance with the invention, a portion of the bandwidth of TDM bus 204 is pre-assigned to all the packet data. In particular, this packet-dedicated portion of the bandwidth is referred to herein as a "multiple-access packet channel," which is shared among at least two of the packet application modules. This is in contrast to allocating a fixed fraction of the TDM bandwidth to each packet application module. The "multiple-access packet channel" provides a single channel for communicating all packet data to, or from, NAM 205. In effect, this "multiple-access packet channel" resembles a packet "local area network" (LAN) in many respects, except that the bandwidth of the "multiple-access packet channel"

5,719,858

5

is closely matched (or equal) to that allocated for packet traffic across network facility 206. In this case, each packet application module must contend for the bandwidth of the "multiple-access packet channel" with the other packet application modules. (Although, the nature of a TDM bus prevents any packet application module from using more bandwidth than is available from the network, this approach makes the entire network bandwidth allocated to packet data dynamically available to each packet application module as a channel for the transport of packet data.)

NAM 205 communicates to all application modules and controls time-slot allocation among the synchronous modules and the packet modules. The control of time-slot allocation can be performed either over TDM bus 204 or another bus (not shown). In this example, it is assumed that a portion of the TDM bus bandwidth is allocated to control and status information, as known in the art. In accordance with the principles of the invention, the time-slots allocated by NAM 205 to the packet application modules are the "multiple-access packet channel."

In this example, it is assumed that every time-frame is comprised of a plurality of 64 Kbit (DS0) channels, and it is assumed that the "multiple-access packet channel" is any grouping of these DS0 channels. Although not a requirement, it is also assumed that the "multiple-access packet channel" is composed of a subset of contiguous DS0 channels available on the bus. In this case, this corresponds to time-slots 1, 2, 3, 4, 5, and 6, of every frame. The remaining time-slots are allocated to synchronous data and control/status information. The allocation of the above-mentioned six time-slots yields a "multiple-access packet channel" with a bandwidth of 384 Khz. The number of DS0 channels allocated on the inbound and outbound data highways are assumed to be the same, so that equal bandwidth is assured in both directions. (As described above, it is assumed that TDM bus 204-i and TDM bus 204-o are symmetrical, i.e., time-slots 1 through 6 are used for the "multiple-access packet channel" for inbound and outbound traffic.) Finally, it is assumed that this 384 Khz bandwidth is equal to the bandwidth allocated over the network facility for the same packet traffic, i.e., NAM 205 is not required to buffer the packet data.

In accordance with a feature of the invention, it is assumed that packet transmission is utilizing HDLC, which is a bit-oriented layer 2 protocol that produces variable bit length packets. This allows a packet to be spread across time-slots and multiple TDM frames. The "multiple-access packet channel" in this case can be considered a continuous sequence of bits with no framing boundaries other than those of the protocol. As such, the starting point of a packet may lie anywhere in the "multiple-access packet channel." An illustrative sequence of frames is shown in FIG. 5. Frame 1 includes time slots 1 through N, where, as mentioned above, each time-slot represents a 64 Kbit DS0 channel. In this example, it is assumed that N is equal to 64 time-slots. Each frame repeats every 125 micro-seconds and each time-slot is further broken down into a sequence of 8 bits as shown in FIG. 5. Each bit is hereafter referred to as a "bit time-slot." It should be noted that the term "time-slot" refers to a collection of "bit time-slots." Time-slots 1-6 represent the "multiple-access packet channel." As described above, since HDLC is a bit-oriented layer 2 protocol, this allows the packet data to begin and end anywhere in a time-slot. Consequently, each packet is mapped into a plurality of "bit time-slots" within time-slots 1 through 6 on TDM bus 204, rather than the DS0 channels representing each time-slot as a whole. This is illustrated in FIG. 5, where packet 50 begins

6

with "bit time-slot" 7 in time-slot 4, of frame 1, and ends with "bit time-slot 3" in time-slot 2 of frame 4. As illustrated by the starting and ending time of packet 50, of FIG. 5, these bit intervals do not have to conform to time slot boundaries into which the packets are mapped and inserted.

As described above, each packet application module must contend for the "multiple-access packet channel." If a packet application module "grabs" the "multiple-access packet channel" that packet application module then transmits using the full 384 Khz of bandwidth. However, if a packet application module cannot "grab" the "multiple-access packet channel," then that packet application module must queue, or buffer, the packet. As a result, the flow control is now distributed among the packet application modules.

The properties desired for packet transport on the "multiple-access packet channel" are high bandwidth efficiency and deterministic fairness. Bandwidth efficiency is a measure of how closely the packet utilization of the available (fixed) TDM bandwidth matches the offered packet load. It also measures efficiency of the channel as the offered packet load exceeds the available bandwidth. Fairness describes how a chosen priority scheme affects the comparative delay of packets through the system under bandwidth contention with multiple packet sources.

Any method used for implementing a "multiple-access packet channel" should be designed to achieve as close to 100% bandwidth efficiency as possible with no negative throughput effects due to congestion. The method should also have no inherent priority bias among the packet sources so that priorities may be enforced selectively if needed, preferably in software on the packet application modules. These properties are only applicable to the outbound (toward the network) direction, where multiple packet sources are contending for a fixed network pipe. Inbound packet data is already multiplexed as a single packet stream within the network channel. Since the TDM bus 204 is full duplex, as represented by separate data highways TDM bus 204-i and TDM bus 204-o, NAU 200 may have an asymmetric access protocol. That is, the access protocol for the inbound direction can be different from the access protocol described below for the outbound direction. In the context of this invention, it is assumed that all modules are continually listening to the TDM in-bound bus to pull off data addressed to them. In the case of the packet application modules, it is assumed that each packet application module is monitoring, or listening to, the "multiple-access packet channel" for packets that have addresses associated with that packet application module. A packet application module listens for its header, e.g., virtual address, if it is not their packet, it is just dropped. The remainder of this description will focus on outbound packet traffic.

In accordance with the invention, a Time-division Multiple Access with Collision Avoidance (TDMA/CA) scheme is used for the outbound direction to regulate access to the "multiple-access packet channel." In particular, the synchronous property of TDM bus 204 provides the means to implement a slotted-access method to avoid collisions. This slotted-access method enables each packet application module to contend for the "multiple-access packet channel," and avoid packet collisions, in a fair and efficient manner. In this embodiment, the slotted-access method is implemented by packet/TDM interface 310 of each packet application module.

Before describing it in detail, an overview of this slotted-access method is described. In the slotted-access method, each packet application module, in rotation order, is given an

5,719,858

7

"access window" of time, corresponding to a "bit time-slot" on the TDM bus, to either capture the "multiple-access packet channel" for transmission, or defer and allow the "access window" to advance to the next packet application module in order. Once granted access, that packet application module has sole access to the "multiple-access packet channel" for a period of time, referred to herein as the "access period." During this "access period," a packet application module sends at least one HDLC frame of queued packet traffic toward the network, and the "access window" is frozen at the current packet application module and does not advance until that packet application module releases the "multiple-access packet channel." Under some conditions however, a packet application module may not begin transmitting packet data as soon as it has captured the "multiple-access packet channel." In particular, whenever the previously transmitting packet application module is still transmitting a packet or closing flag, the next packet application module must wait until completion before transmitting its first packet. Since a rotational arbitration scheme on the bus may take several "bit time-slots," or clock, intervals, this mechanism allows the arbitration to overlap an active packet transmission, avoiding idle time on the bus and increasing bandwidth utilization.

FIG. 6 shows an illustrative slotted-access method. There is an implied numbering of the packet application modules and "bit time-slots." This implied numbering is hereafter referred to as an "ID number." Generally speaking, a packet application module can attempt to access TDM bus 204-o only when the ID number of the "bit time-slot" matches the ID number of the packet application module. In a system with N packet application modules, each "bit time-slot" of the "multiple-access packet channel" is counted in a repeating sequence from 0 to N-1 starting at some arbitrary "bit time-slot" by each packet application module. In other words, each packet application module only counts those "bit time-slots" assigned to the "multiple-access packet channel." The value of the count is the ID number for the "bit time-slot." All packet application modules are synchronized with this counting sequence and are aware of the ID number of each "bit time-slot" at all times. As a result, the ID number need not be present on the bus. As noted earlier, each packet application module knows, a priori, the time-slots associated with the "multiple-access packet channel." In combination with this, each packet application module is configured with a unique ID number, which can be determined in any number of ways. For example, each module, or circuit board, can have an address associated either via a "software-controlled" configuration, e.g., a system administrator literally assigns addresses to the various modules; or by a hardware setting that specifically associates a particular address with a particular position in the system, e.g., what slot the circuit board is plugged into. Note, this counting of "bit time-slots" may span more than one frame. Since N may be any integer there is no relationship between this ID numbering and the position of bits in the TDM frame.

For example, if there are seven packet application modules, and assuming for the moment that none of the seven packet application modules wanted access to the TDM bus, this counting would look like that shown in FIG. 7. FIG. 7 is similar to FIG. 5, except packet 50 has been removed, i.e., there is no transmission of a packet and only time-slots 1 and 2 of frame 1 are shown. It is assumed that each packet application module counts "bit time-slots" of the "multiple-access packet channel" beginning with "bit time-slot" 1 of frame 1, which is associated with the count value of 0. Each packet application module waits for its ID number to equal

8

the count, or ID number, of the "bit time-slot" to attempt to access TDM bus 204-o. For example, the packet application module associated with ID 0, can attempt access only upon the value of the count equaling 0, which, in this example, occurs in "bit time-slot" 1 of time-slot 1 of frame 1, "bit time-slot" 8 of time-slot 1 of frame 1, "bit time-slot" 7 of time-slot 2 of frame 1 etc.

To implement this slotted-access method two additional signals are bussed between packet application modules. These signals are "packet request" (PREQ) and "packet hold" (PHOLD). It is assumed these signals are bussed among the packet application modules as simply "open collector" as known in the art which allows them to be logically "OR"ed. Referring back to FIG. 6, a sequence of "bit time-slots" is shown. This sequence of bit time-slots only represents those "bit time-slots" assigned to the "multiple-access packet channel." The top row of FIG. 6 is simply a sequence of bit time reference points. The second row of FIG. 6 is the ID number of the "bit time-slot" of the "multiple-access packet channel," i.e., the value of the count. The next two rows represent the state of the PREQ and PHOLD busses. The final row simply represents packet data.

Beginning at time t2, the "bit time-slot" ID number is equal to 0. In order for a packet application module to transmit a packet on TDM bus 204-o, the packet application module must assert the PREQ signal whenever the ID numbers of the "bit time-slot" and the packet application module match. Upon receiving the PREQ signal, each packet application module halts the counting. The packet application module that asserted the PREQ signal becomes the next in line to begin transmission. For example, when the packet application module associated with the ID number of 2 wants to transmit, it must wait until time t4, when the ID numbers of the "bit time-slot" and the packet application module match, to assert the PREQ signal. However, the presence of a PHOLD signal driven by the currently transmitting packet application module delays the access of packet application module 2 to TDM bus 204-o until PHOLD is withdrawn at time t5. (The question marks illustrated in FIG. 6 simply represent that another, unidentified packet application module is currently transmitting.) At the next "bit time-slot," t6, packet application module 2 concatenates its packet stream with that of the previous transmission.

Although it could occur at any point, it is assumed that packet application module 2 drops the PREQ signal at the end of its packet transmission (described below) to allow the counting of time-slots by all packet application modules to again advance. At the same time, packet application module 2 asserts the PHOLD signal, which prevents the next packet application module from beginning its transmission until packet application module 2 has completed its closing flag sequence. These events occur at time t9. Subsequently, packet application module 6 raises its PREQ signal at time t13, signaling its readiness to send at least one packet, while packet application module 2 is still transmitting its closing flag. At time t18, packet application module 2 completes its packet transmission and drops PHOLD. At time t19, packet application module 6 begins sending its packet data.

It should be noted that the HDLC flags are a byte in length. Consequently, transmission of an HDLC inter-frame flag must end first before another packet application module can get the TDM bus 204-o to insert data. In addition, and in accordance with HDLC, the last packet application module may continue inserting the inter-frame flags, and asserting the PHOLD signal until the PREQ signal is asserted, by

5,719,858

9

itself or another module. Once the PREQ signal is asserted, TDM/packet interface 310 drops the PHOLD signal only when the insertion of the current flag is finished.

As mentioned above, the point at which an actively transmitting packet application module asserts the PHOLD signal and lowers the PREQ signal is arbitrary. The earlier this occurs, the higher the probability that the next packet application module will be found by the time the first closing flag sequence is sent. This increases the efficiency of the "multiple-access packet channel" because there is no waiting for data transmission to finish to start arbitration again, i.e., no time is lost for contention (a packet application module is always ready to go). On the other hand, waiting until near the end of the transmission allows more timely packet queue information to be used in the arbitration. A balance between these two extremes could be implemented.

A method illustrating this slotted-access technique for use in the packet application module of FIG. 4 is shown in FIG. 8. The steps shown in FIG. 8 have been divided into different "subroutines" for simplicity, i.e., an "initialize subroutine" (FIG. 8A), "request to transmit a packet subroutine" (FIG. 8B), and a "wait for PHOLD" subroutine (FIG. 8C). In step 505, packet application module 215-n is initialized, e.g., via a power-up sequence, and determines its ID number, e.g., via a hardware strap or software configuration. During this initialization, NAM 205 of FIG. 3 allocates the time-slots associated with the "multiple-access packet channel." Packet/TDM interface 310 uses the assignment information from NAM 205 as the synchronizing signal to begin counting "bit time-slots" of the "multiple access packet channel" in step 510 (provided that the PREQ signal is not asserted). In this example, it is assumed that packet/TDM interface 310 includes a counter to count each "bit time-slot" of TDM bus 204-o that is associated with the "multiple-access packet channel." The counter included within packet/TDM interface 310 runs continuously and is controlled by the PREQ signal, i.e., if PREQ is asserted—no counting takes place and the count holds at the last value.

When there is a packet to transmit, packet interface processor 305 begins to fill buffer 315 provided buffer 315 is not full. In particular, referring briefly back to FIG. 4, it can be seen that a BUFFER FULL signal is provided by buffer 315 to packet interface processor 305. This signal alerts packet interface processor 305 if buffer 315 is full. As a result, packet interface processor 305 fills buffer 315 when packet data is available to transmit only if the BUFFER FULL signal is not active. Once the BUFFER FULL signal is active, packet interface processor 305 could, if necessary, perform flow-control, if possible, with the associated packet endpoint like a router, or simply begin dropping packets. Assuming buffer 315 is initially empty, as packet interface processor 305 begins to fill up buffer 315, the BUFFER NOT EMPTY signal is asserted for TDM/packet interface 310 in step 520, via line 321.

Upon receiving the BUFFER NOT EMPTY signal, TDM/packet interface 310 waits for the "access window" in step 525. Once the "access window" matches, i.e., the ID number of packet application module 215-n matches the current value of the counter, i.e., the "bit time-slot" ID number, TDM/packet interface 310 asserts the PREQ signal in step 530.

In step 535, TDM/packet interface 310 waits until PHOLD is no longer asserted before transmitting the packet from buffer 315 onto TDM bus 204-o. Upon nearing the end of the packet transmission (which can be determined simply by knowing the length of the packet from the packet header

10

information), TDM/packet interface 310 drops the PREQ signal and asserts the PHOLD signal in step 540. As mentioned above, the last packet application module to grab the "multiple-access packet channel" continues inserting flags and asserting PHOLD until PREQ is asserted, by itself or another module. Once PREQ is again asserted, TDM/packet interface 310 drops PHOLD only when the insertion of an flag is finished. At this point the next packet application module then has access to the "multiple-access packet channel."

Within this general method, different approaches may be taken to offset the arbitration fairness. In general, each time a packet application module gains transmission access to the "multiple-access packet channel," it may send any prescribed number of packets. The access method may be designed to limit this number to one, or it may allow the application to empty its packet transmit buffer. If the hardware imposes no limit on the size or number of packets sent during one access period, the application software is then able to implement rules for sending varying amounts of packet traffic across the bus. There is no reason why those rules must be applied the same for every packet application module, and in fact could be different for each port.

For example, the packet application module may send as many packets as can be transmitted within a fixed amount of time. Another possibility is that the amount of packet data transmitted may be a function of the number of packets queued or the time they have been queued. The size of the packets, always known to the software, may also be a factor in determining when to terminate the access period.

The only hardware function required to support such strategies is the combination of the PREQ signal and PHOLD signal. Once the active packet application module is about to fulfill its prescribed transmission requirement, the corresponding packet interface processor either informs the respective TDM/packet interface via a control line (not shown) or the TDM/packet interface continues transmitting until a corresponding BUFFER EMPTY signal (not shown) is received by the TDM/packet interface. Whatever the signal, the TDM/packet interface then stops asserting PREQ to allow another packet application module to gain access to the TDM bus.

As noted above, the above description assumed that the packet/TDM interface was incorporated in a single ASIC bus interface chip. This is feasible if bus speeds and capacities are modest, say 10-20 Mbps maximum throughput, and only one packet channel is required. As a result, the use of an ASIC to implement a bus interface chip to integrate these functions should result in minimal additional cost (in quantity). Further, adding support for a second packet channel may be a similar incremental cost. However, boosting the TDM bus capacity to T3 rates or higher will pose the greatest potential cost increase due to higher bus clock rates and the number of 64K channels. At this level, an ASIC having a much larger gate-count may be needed to provide a large number of timeslot registers, increased buffering, and higher speed DMA channels.

The foregoing merely illustrates the principles of the invention and it will thus be appreciated that those skilled in the art will be able to devise numerous alternative arrangements which, although not explicitly described herein, embody the principles of the invention and are within its spirit and scope.

For example, although the invention is illustrated herein as being implemented with functional building blocks, e.g., a packet interface processor, etc., the functions of any one or

5,719,858

11

more of those building blocks can be carded out using one or more appropriate integrated circuits.

In addition, although illustrated in the context of a T1 network facility, the inventive concept applies to other network facilities as well, e.g., fractional T1, digital data service (DDS), T3, etc. Further, more than one "multiple-access packet channel," can exist. For example, a first plurality of packet application modules may be assigned to a first group of time-slots, e.g., time-slots 1-6, for transmitting packet data, while a second plurality of packet application modules is assigned to a second group of time-slots, e.g., time-slots 7-12, for transmitting packet data. Also, because the bandwidth of a TDM bus may be divided into many separate logical channels, data with different formats and access methods, such as isochronous and packet data, may be combined in the system. Mother variation is to use the time-slots to control contention access, as opposed to the "bit-time slots," described above.

It should also be noted that the inventive concept is applicable to Asynchronous Transfer Mode (ATM) transmission. In particular, ATM cells are handled in a similar manner, although not on the same "multiple-access packet channel" carrying bit-oriented protocols. In the case of ATM cells, the octets which form the cells need to be aligned within DS0 channels. This places an additional constraint on the packet/TDM interface to recognize those boundaries and transmit within them. The "bit time-slots" may still be used in the same manner as describe above for arbitration however. Another property of ATM cells is that they are not delimited by flags. Idle time between cells transporting user data must be filled with null cells. Responsibility for filling voids in the cell stream may be assigned to the packet application modules to perform in a fashion similar to adding flag fills for HDLC protocols described earlier. Alternatively, this could be the responsibility of the NAM on the outbound trunk, especially if the NAM is required to buffer cells in order to map them onto the network link.

What is claimed:

1. Data communications apparatus comprising:
 - a time division multiplexed bus having a bandwidth, where a portion of the bandwidth is allotted to packet data;
 - a plurality of packet data sources coupled to the time-division multiplexed bus that share the allotted bandwidth for transmitting packet data; and
 - a distributed packet manager within each of said packet data sources configured to allocate access to the allotted bandwidth among said packet data sources.
2. The apparatus of claim 1 including a packet arbitration bus comprising a packet request signal and a packet hold signal.
3. The apparatus of claim 2 wherein each packet data source comprises:
 - interface circuitry that inserts packets to the allocated bandwidth of the time-division multiplexed bus and is coupled to the packet arbitration bus; and
 - a packet data processor for providing the packets from a packet endpoint to the interface circuitry.
4. Data communications apparatus comprising:
 - a time-division multiplexed bus having a bandwidth, where a portion of the bandwidth is allocated to packet data;
 - a plurality of packet data sources coupled to the time-division multiplexed bus that share the allocated bandwidth for transmitting packet data;
 - a packet arbitration bus comprising a packet request signal and a packet hold signal;

12

interface circuitry coupled to the packet arbitration bus, said interface circuitry inserting packets to the allocated bandwidth of the time-division multiplexed bus, and performing a counting function by counting time-slots of the allocated bandwidth so long as the packet request signal is not asserted; and

a packet data processor for providing the packets from a packet endpoint to the interface circuitry.

5. The apparatus of claim 4 wherein the time-slots are bit time-slots.

6. The apparatus of claim 4 wherein the interface circuitry inserts the packets into the allocated bandwidth when a value of the counter is equal to a predetermined identification number of the respective packet data source and the packet hold signal is not asserted.

7. Communications apparatus comprising:

a time-division multiplexed bus having a predefined bandwidth;

a plurality of synchronous data sources coupled to the time-division multiplexed bus for communicating synchronous data in a first portion of the predefined bandwidth;

a plurality of packet data sources coupled to the time-division multiplexed bus for communicating packet data in a second portion of the predefined bandwidth, where the plurality of packet data sources share the second portion of the predefined bandwidth for transmitting packet data; and

a distributed packet manager within each of said packet data sources configured to allocate access to the second portion of the predefined bandwidth among said packet data sources.

8. The apparatus of claim 7 further including a network access manager coupled to the time-division-multiplexed bus for communicating the synchronous data and the packet data to at least one network facility.

9. Communications apparatus comprising:

a time-division multiplexed bus having a predefined bandwidth;

a plurality of synchronous data sources coupled to the time-division multiplexed bus for communicating synchronous data in a first portion of the predefined bandwidth; and

a plurality of packet data sources coupled to the time-division multiplexed bus for communicating packet data in a second portion of the predefined bandwidth, where the plurality of packet data sources share the second portion of the predefined bandwidth for transmitting packet data, the second portion of the predefined bandwidth being shared in such a way that only one of the plurality of packet data sources accesses the second portion of the predefined bandwidth at a time.

10. The apparatus of claim 7 wherein each one of the plurality of packet data sources includes interface circuitry to the time-division multiplexed bus for synchronizing packet data to the time-division multiplexed bus.

11. Communications apparatus comprising:

a time-division multiplexed bus having a predefined bandwidth;

a plurality of synchronous data sources coupled to the time-division multiplexed bus for communicating synchronous data in a first portion of the predefined bandwidth;

a plurality of packet data sources coupled to the time-division multiplexed bus for communicating packet

5,719,858

13

data in a second portion of the predefined bandwidth, where the plurality of packet data sources share the second portion of the predefined bandwidth for transmitting packet data, wherein each one of the plurality of packet data sources includes interface circuitry to the time-division multiplexed bus for synchronizing packet data to the time-division multiplexed bus, and the interface circuitry includes a counter for counting time-slots representing the second portion of the predefined bandwidth.

12. The apparatus of claim 11 wherein the time-slots are bit time-slots.

13. The apparatus of claim 11 wherein the interface circuitry inserts packets into the second portion of the predefined bandwidth when a value of the counter equals a predetermined identification number of a respective packet data source and a hold signal is not being asserted by the interface circuitry of the remaining packet data sources.

14. The apparatus of claim 11 wherein the counter is inhibited from counting when a packet request signal is asserted by interface circuitry from any packet data source.

15. A method for use in a data communications apparatus for transmitting packet data on a time-division multiplexed bus, the method comprising the steps of:

coupling a plurality of packet data sources to the time-division multiplexed bus;

allocating a portion of the bandwidth of the time-division multiplexed bus to the plurality of packet data sources in such a way that the allocated portion is shared among the plurality of packet data sources;

transmitting packet data from the plurality of packet data sources on the allocated portion of the bandwidth; and controlling access by said packet data sources to the allocated portion of the bandwidth via a distributed packet manager within each of said packet data sources.

16. A method for use in a data communications apparatus for transmitting packet data on a time-division multiplexed bus, the method comprising the steps of:

coupling a plurality of packet data sources to the time-division multiplexed bus;

allocating a portion of the bandwidth of the time-division multiplexed bus to the plurality of packet data sources in such a way that the allocated portion is shared among the plurality of packet data sources;

assigning to each one of the plurality of packet data sources an identification number;

counting, in each one of the plurality of packet data sources, a sequence of time-slots representing the allocated portion of the bandwidth; and

transmitting packet data from one of the plurality of packet data sources on the allocated portion of the bandwidth only when the value of the count matches the respective identification number assigned to that packet data source.

17. The method of claim 16, wherein said step of counting a sequence of time-slots representing the allocated portion of the bandwidth is enabled only when none of the plurality of packet data sources assert a packet request signal.

18. The method of claim 16, wherein said step of transmitting packet data from one of the plurality of packet data sources is enabled only if a packet hold signal is not asserted by any other one of the plurality of packet data sources.

19. The method of claim 16, wherein said step of counting a sequence of time-slots representing the allocated portion of the bandwidth further comprises counting bit time-slots.

20. A method for transmitting packet data on a time-division multiplexed bus in data communications equipment, the method comprising the steps of:

14

allocating a portion of the bandwidth of the time-division multiplexed bus as a multiple-access packet channel; coupling a plurality of packet data sources to the time-division multiplexed bus;

controlling the access by said packet data sources to the allocated portion of the bandwidth via a distributed packet manager within each of said packet data sources;

transmitting packet data from the one of the plurality of packet data sources having access to the multiple-access packet channel.

21. A method for transmitting packet data on a time-division multiplexed bus in data communications equipment, the method comprising the steps of:

allocating a portion of the bandwidth of the time-division multiplexed bus as a multiple-access packet channel;

coupling a plurality of packet data sources to the time-division multiplexed bus;

assigning a unique identification number to each one of the plurality of packet data sources;

counting, in each one of the plurality of packet data sources, a sequence of time-slots representing the multiple-access packet channel;

granting access to the multiple-access packet channel to that one of the plurality of packet data sources that requests access when the value of the count matches the respective identification number of that packet data source; and

transmitting packet data from the one of the plurality of packet data sources having access to the multiple-access packet channel.

22. The method of claim 21 wherein the counting step is enabled only when none of the plurality of packet data sources assert a packet request signal.

23. The method of claim 21 wherein the granting step further includes the step of asserting a packet request signal by the packet data source granted access.

24. The method of claim 21 wherein the counting step counts a sequence of bit time-slots.

25. A method for transmitting packet data on a time-division multiplexed bus for use in data communications equipment, the method comprising the steps of:

allocating a portion of the bandwidth of the time-division multiplexed bus as a multiple-access packet channel;

coupling a plurality of packet data sources to the time-division multiplexed bus;

arbitrating between the plurality of packet data sources to grant access to the multiple-access packet channel to one of the plurality of packet data sources at a time; and transmitting packet data from the one of the plurality of packet data sources having access in the arbitrating step only if a packet hold signal is not asserted by another one of the plurality of packet data sources.

26. The method of claim 20 further comprising the step of coupling a network access module to the time-division multiplexed bus for receiving the packet data for transmission over a network facility.

27. A data communications apparatus, comprising:
a time-division multiplexed bus with a bandwidth which is split into at least two portions, one of the portions being allocated for packet data;

a plurality of packet data sources coupled to said time-division multiplexed bus, the packet data sources transmitting packet data on said portion allocated for packet data; and

5,719,858

15

means for arbitrating access to the portion of the time-division multiplexed bus allocated for packet data among the plurality of packet data sources.

28. The apparatus of claim 27, wherein said means for arbitrating access to the portion of the time-division multiplexed bus allocated for packet data further comprises a packet arbitration bus having a packet request signal and a packet hold signal, the arbitration bus being coupled to each of the packet data sources, and the packet data sources including means for generating the packet request signal and the packet hold signal.

29. The apparatus of claim 28, wherein said means for arbitrating access to the portion of the time-division multiplexed bus allocated for packet data further comprises a means within each packet data source for counting the time-slots of the portion of the time-division multiplexed bus.

30. The apparatus of claim 29, wherein said means for arbitrating access to the portion of the time-division multiplexed bus allocated for packet data further comprises the communication of packet data by one of the packet data

16

sources when the packet hold signal is not asserted and the value of said counting means is equal to a predetermined identification number assigned to said packet data source.

31. The apparatus of claim 29, wherein said means for arbitrating access to the portion of the time-division multiplexed bus allocated for packet data further comprises means for generating the packet request signal by a packet data source wishing to gain access to the time-division multiplexed bus when the value of said counting means is equal to a predetermined identification number assigned to said packet data source, said counting means being stopped by said packet request signal.

32. The apparatus of claim 29, wherein said means for arbitrating access to the portion of the time-division multiplexed bus allocated for packet data further comprises means for releasing said packet request signal and for asserting said packet hold signal by a packet data source while said packet data source is transmitting packet data.

* * * * *

United States Patent [19]

King

[11] Patent Number: 4,937,819

[45] Date of Patent: Jun. 26, 1990

[54] TIME ORTHOGONAL MULTIPLE VIRTUAL DCE FOR USE IN ANALOG AND DIGITAL NETWORKS

[75] Inventor: Joseph B. King, St. Petersburg, Fla.

[73] Assignee: A.T. & T. Paradyne, Largo, Fla.

[21] Appl. No.: 249,450

[22] Filed: Sep. 26, 1988

[51] Int. Cl.³ H04J 3/16

[52] U.S. Cl. 370/95.3; 370/95.1; 370/85.7; 370/104.1

[58] Field of Search 370/96, 90, 85, 86, 370/95, 79, 82, 94, 104, 95.1, 95.2, 95.3, 85.7, 104.1, 94.1; 375/12, 13, 14; 340/825.5, 825.51

[56] References Cited

U.S. PATENT DOCUMENTS

4,398,289	8/1983	Schoute	370/95
4,489,416	12/1984	Stuart	375/13
4,606,023	8/1986	Dragoo	370/104
4,644,534	2/1987	Sperlich	370/95
4,653,049	3/1987	Shinmyo	370/95
4,669,090	5/1987	Betts et al.	375/13
4,694,453	9/1987	Kobayashi et al.	370/85
4,726,017	2/1988	Krum et al.	370/85
4,742,512	5/1988	Akashi et al.	370/96
4,748,621	5/1988	Ballance et al.	370/95
4,757,502	7/1988	Meuriche et al.	370/104

4,759,016	7/1988	Otsuka	370/95
4,797,878	1/1989	Armstrong	370/96
4,800,560	1/1989	Aoki et al.	370/104
4,807,259	2/1989	Yamanaka et al.	370/85

OTHER PUBLICATIONS

Mischa Schwartz, "Communication Systems & Techniques", 1966, p. 173.

A. Bruce Carlson, "Communication Systems: An introduction to Signals & Noise in Electrical Communication", 1968, p. 377.

James Martin, "Satellite Communications", 1978, pp. 276-278.

Primary Examiner—Douglas W. Olms

Assistant Examiner—Alpus H. Hsu

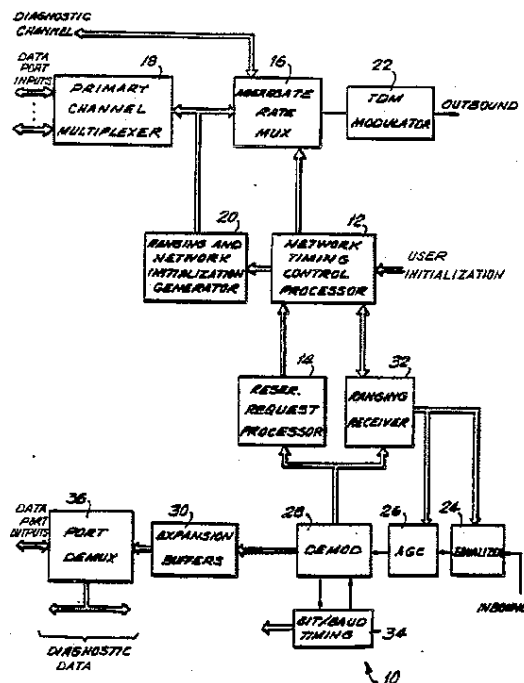
Attorney, Agent, or Firm—Kane, Dalsimer, Sullivan,

Kurucz, Levy, Eisele and Richard

[57] ABSTRACT

Apparatus and method for time division multiple access in a multidrop system with multiple host applications. Employing half-duplex polled protocols is disclosed. Ranging with respect to time is used so as to reduce guard time between successive transmissions thereby increasing system efficiency. Host applications can request extra time slots for long messages via a request bit within the message format.

15 Claims, 9 Drawing Sheets

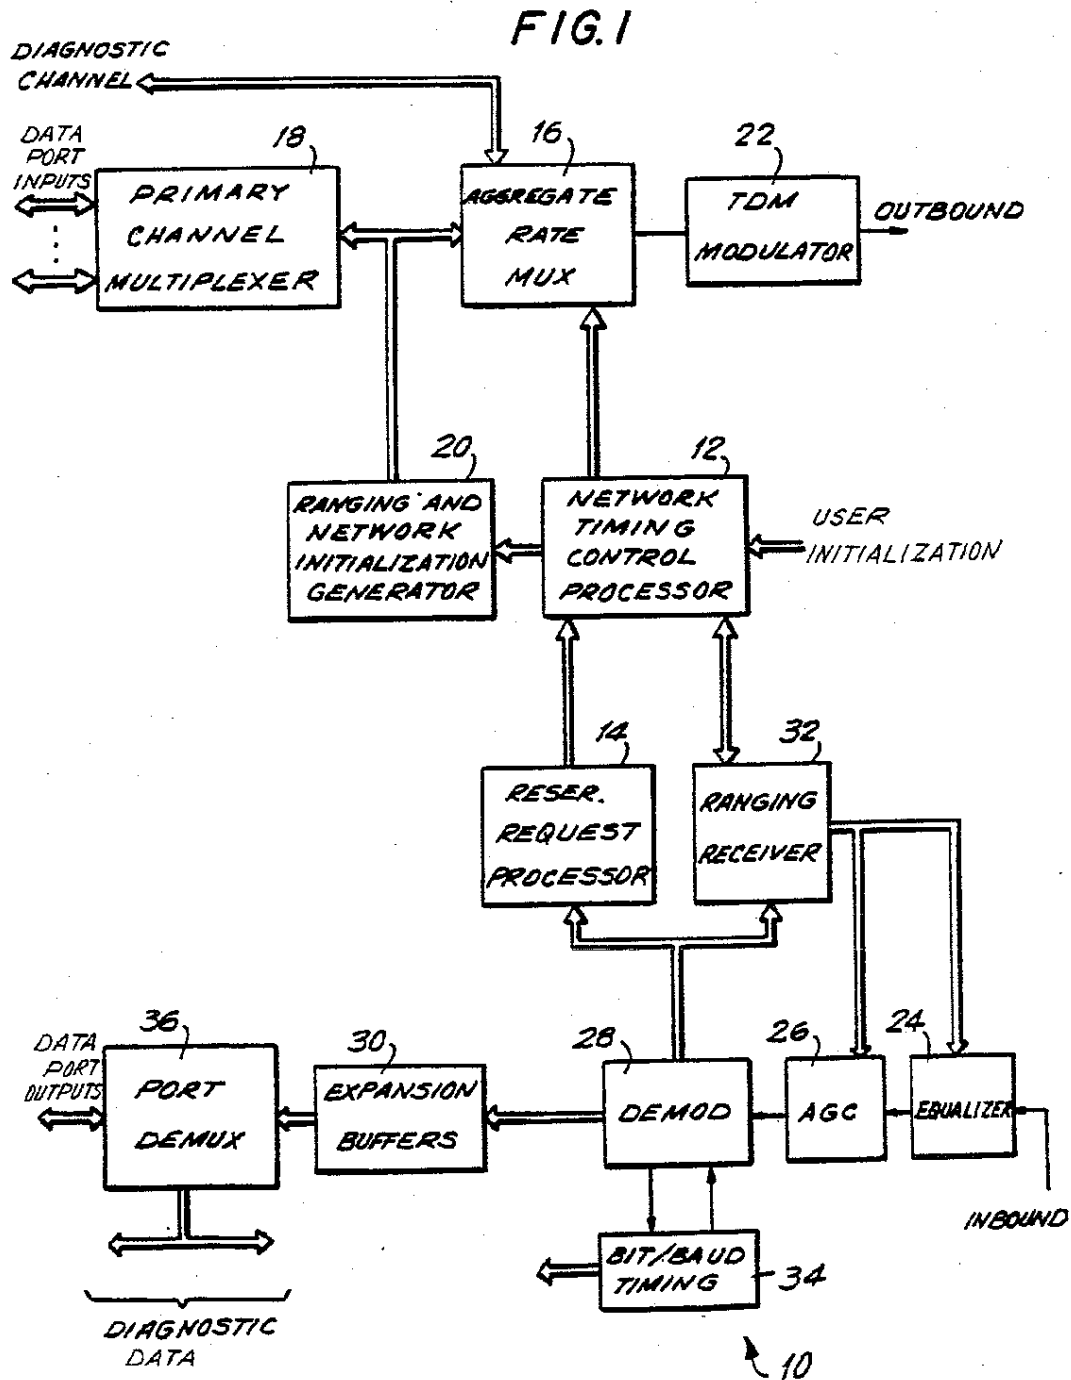



U.S. Patent

Jun. 26, 1990

Sheet 1 of 9

4,937,819





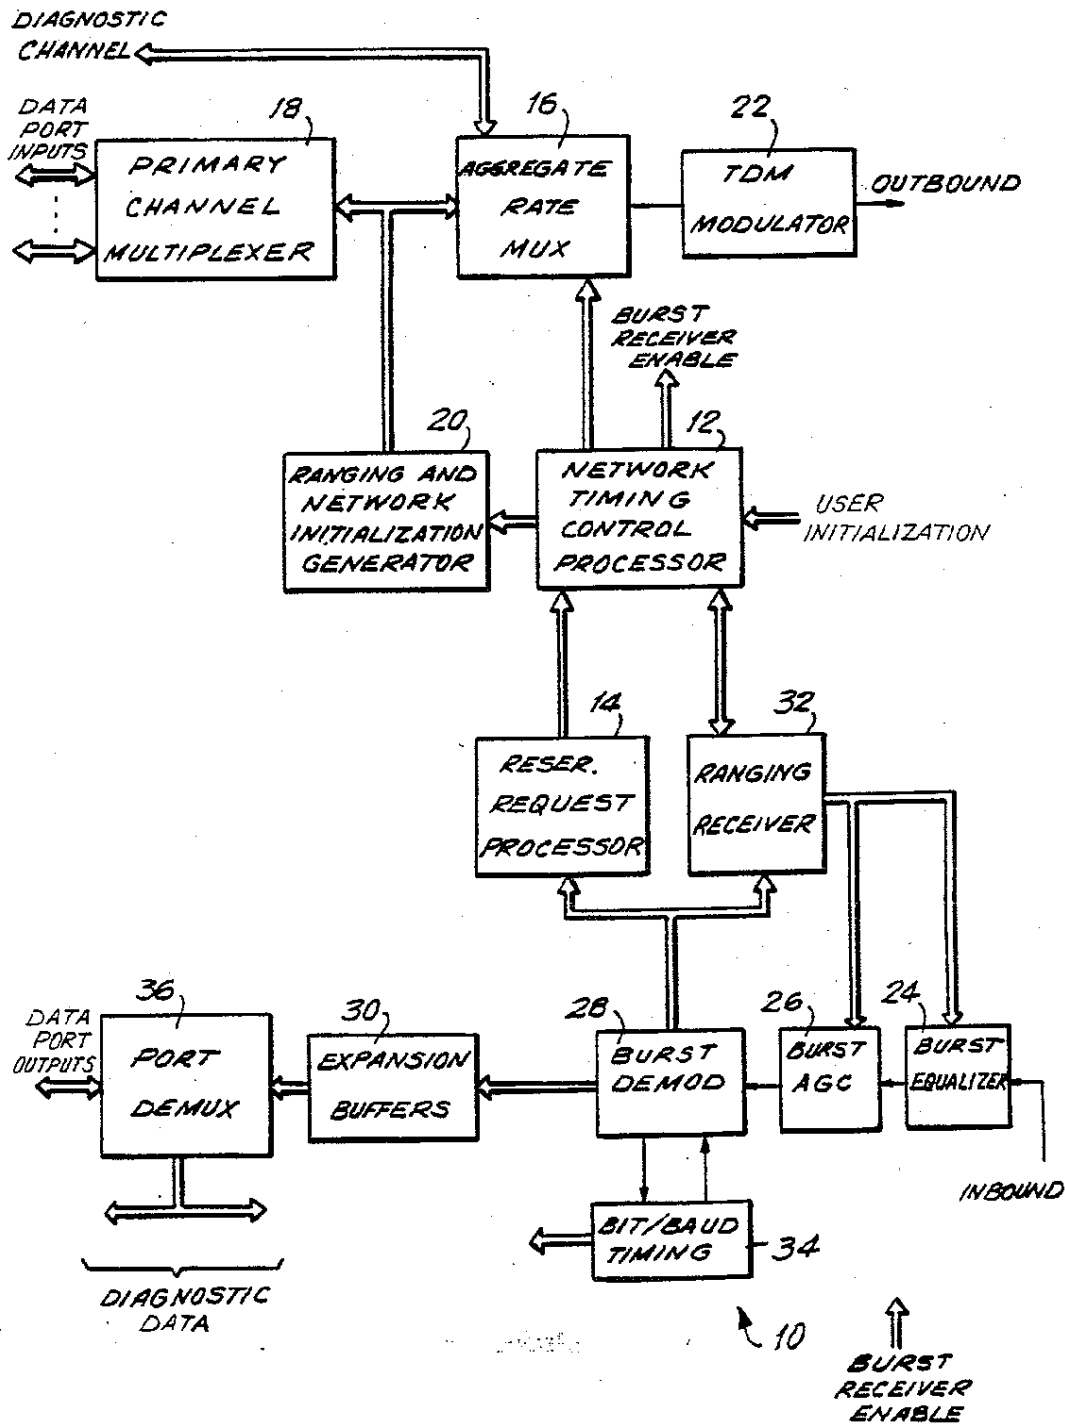
U.S. Patent

Jun. 26, 1990

Sheet 3 of 9

4,937,819

FIG. 3




U.S. Patent

Jun. 26, 1990

Sheet 4 of 9

4,937,819

FIG. 4



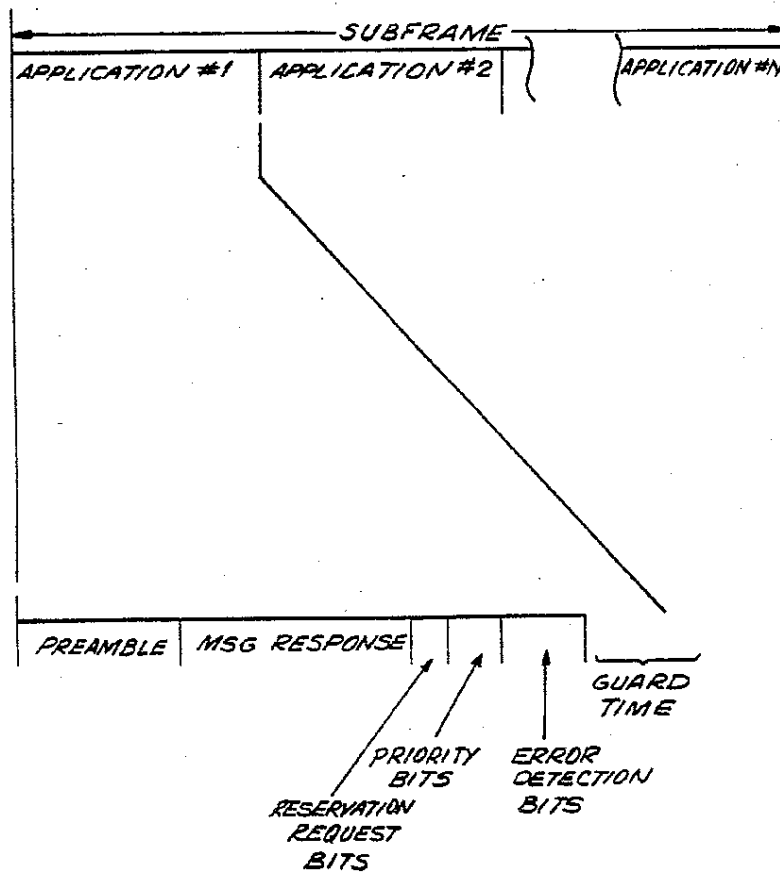
U.S. Patent

Jun. 26, 1990

Sheet 5 of 9

4,937,819

FIG. 5



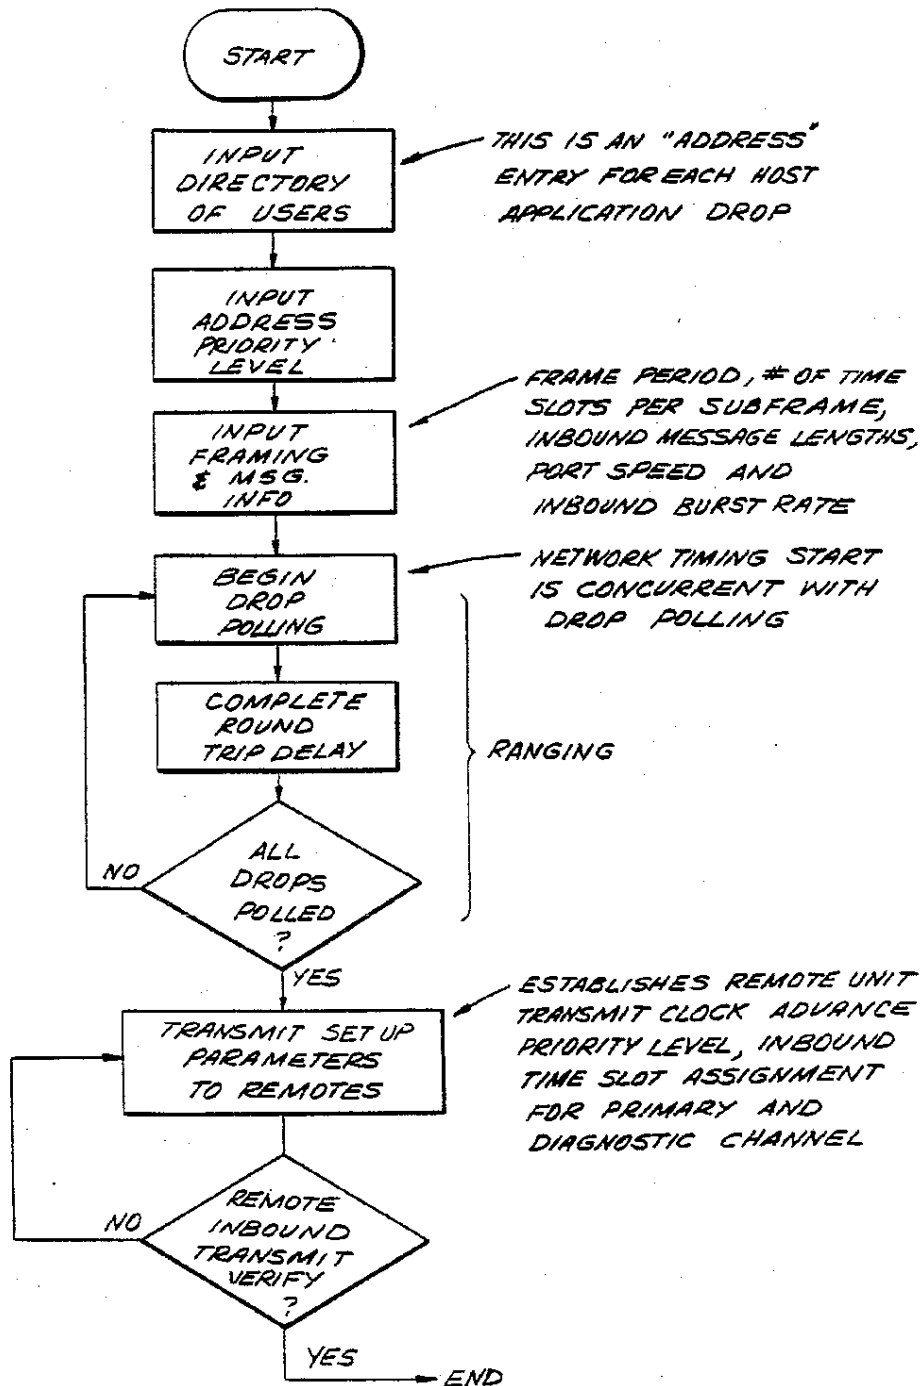
U.S. Patent

Jun. 26, 1990

Sheet 6 of 9

4,937,819

FIG. 6



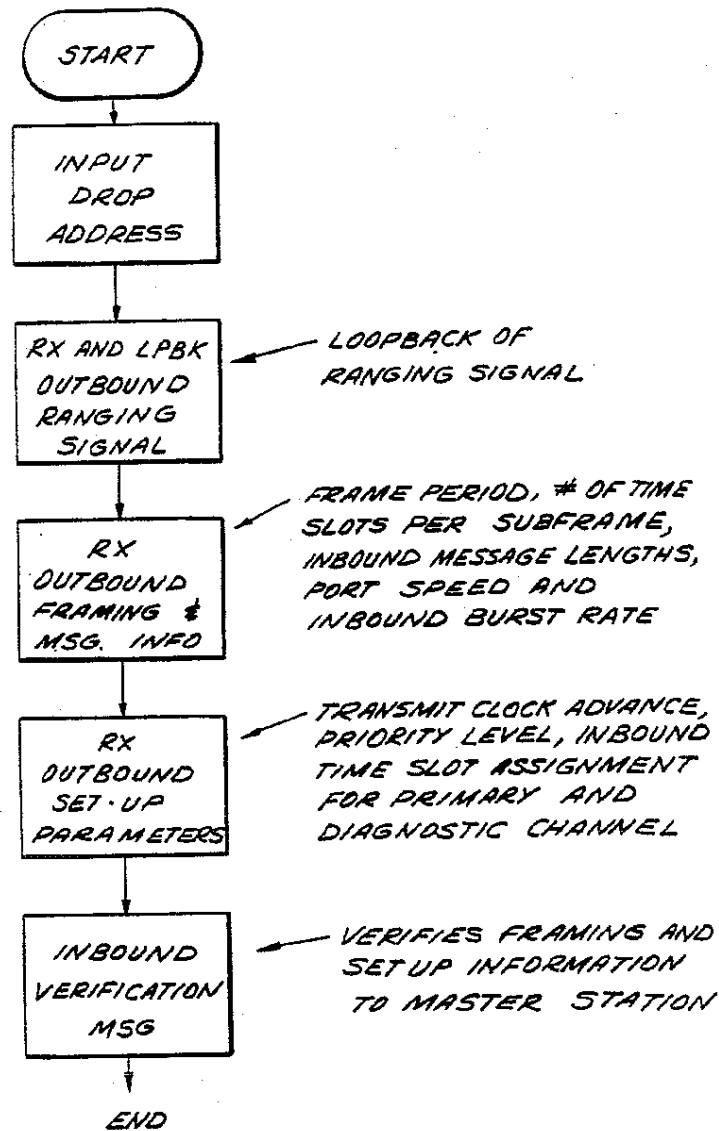
U.S. Patent

Jun. 26, 1990

Sheet 7 of 9

4,937,819

FIG. 7



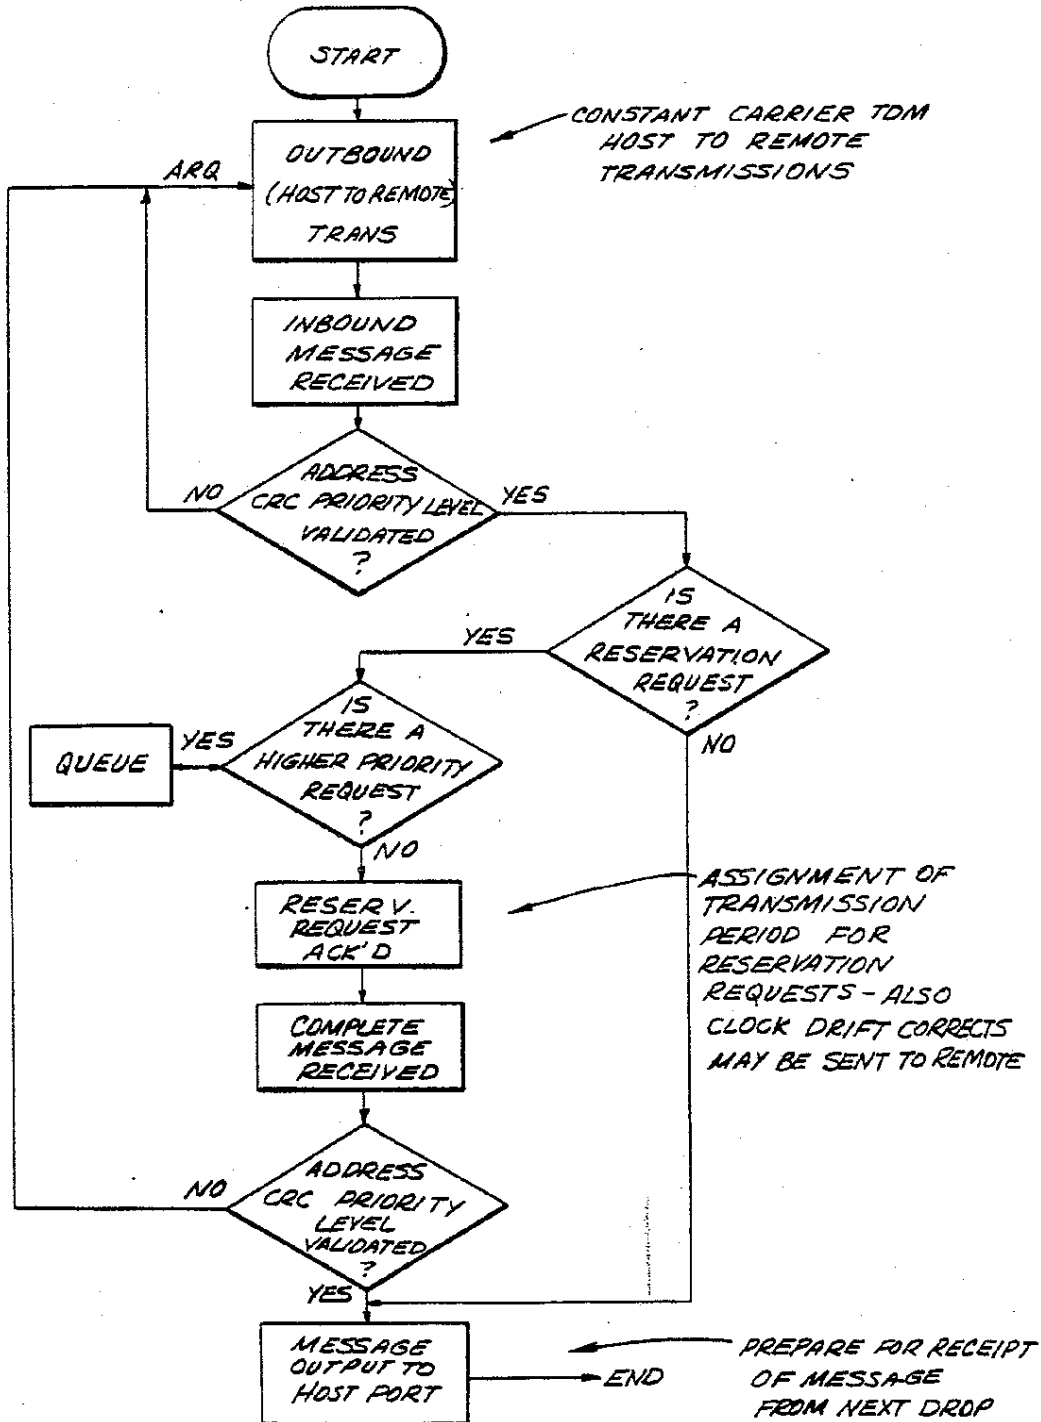
U.S. Patent

Jun. 26, 1990

Sheet 8 of 9

4,937,819

FIG. 8



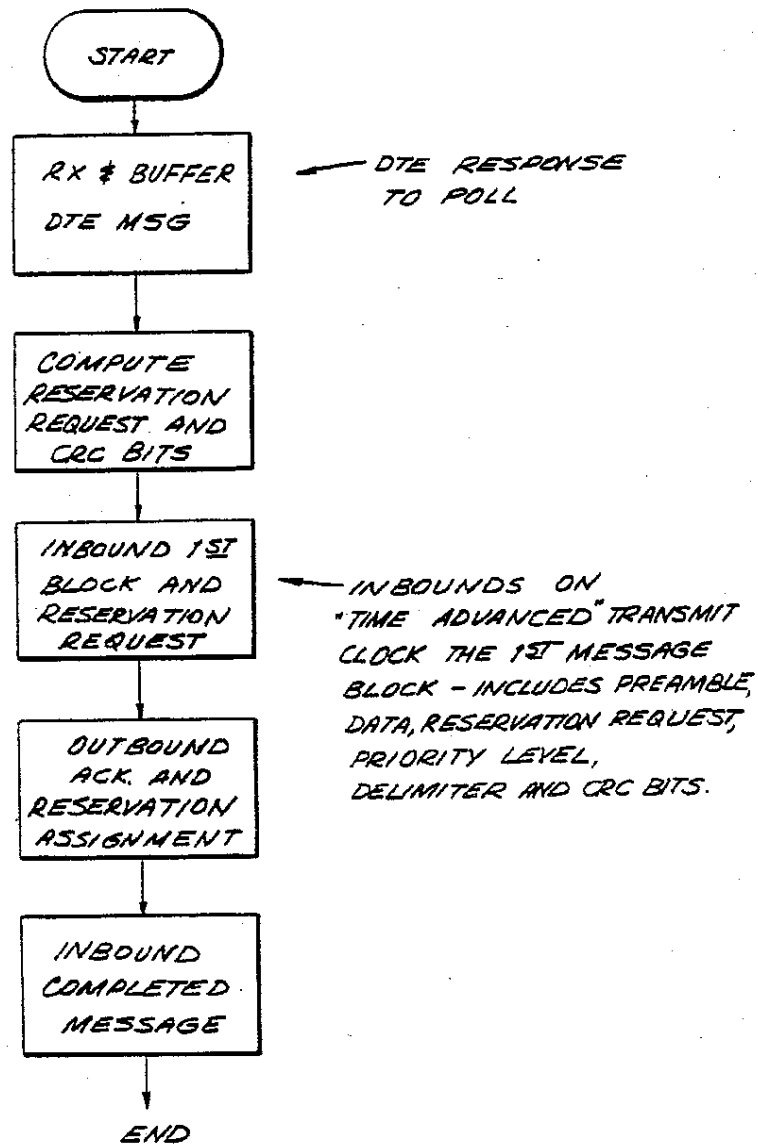
U.S. Patent

Jun. 26, 1990

Sheet 9 of 9

4,937,819

FIG. 9



4,937,819

1

TIME ORTHOGONAL MULTIPLE VIRTUAL DCE FOR USE IN ANALOG AND DIGITAL NETWORKS

BACKGROUND OF INVENTION

1. Field of Invention

This invention relates to an apparatus and method for a master unit in a multidrop network to communicate to and from a plurality of remote units, using a plurality of host applications using half duplex polled protocols, through the use of time division multiple access techniques.

2. Scope of the Prior Art

In the prior art, in order to run multiple host applications to multiple modems in a multidrop network, it is common to use a single network channel for each application, thereby effectively resulting in a number of networks rather than a single network. Further, such an arrangement is clearly an inefficient use of leased lines and other equipment.

A common solution to this deficiency of the prior art is to use a single line with frequency division multiplexing. That is, a number of orthogonal carrier frequencies, one for each application, are transmitted over a single line to a plurality of remote units in a multidrop network. However, with such apparatus, the non-linearities of the communications line (most frequently, a telephone line), interfere with the co-existence of multiple carrier frequencies. This interference includes intermodulation, cross-modulation and spillover between and among channels. Furthermore, a strong signal on one carrier frequency could suppress a weak signal on another carrier frequency.

Due to this interference, time-consuming engineering adjustment is required to install and maintain such a system. Such interference increases with an increasing number of co-existing carrier frequencies, thereby limiting the number of carrier frequencies which could be practically carried on a single line. For practical applications, no more than three carrier frequencies can be carried simultaneously on a telephone line.

OBJECTS AND SUMMARY OF THE INVENTION

It is therefore an object of this invention to provide a method and apparatus for allowing a single multidrop network to run multiple applications to multiple remote units.

It is therefore a further object of this invention to provide such a method and apparatus without a fundamental limitation on the number of applications which can be implemented.

It is therefore a further object of this invention to reduce the amount of engineering adjustment needed to install and maintain such a method and apparatus.

The present method and apparatus permits multiple multidrop networks (such as Dataphone Digital Service for digital applications or conventional telephone company lines for analog applications), each serving a distinct half-duplex host polled application, to be replaced by a single multidrop network serving each of said host applications.

The basic features of this method and apparatus are time division multiplexed outbound transmissions from the master to the remote units for data and control; time division multiple access transmissions inbound from the remote units to the master unit; master to remote ranging with respect to transmission time; and priority as-

2

signed reservation request for long poll responses. A channel rate exceeding the aggregate port rate is required in order to transmit effectively all of the information from the remote units and allow for control format messages. All remote units (or "drops") receive messages outbound from the control unit and respond in a unique time period assigned to each host application. Contention between applications is thereby avoided due to the fact that each application is assigned such a unique time period. By ranging or measuring the round-trip transmission or delay time between the master unit and each remote unit, and storing these times in a table so as to accurately synchronize the transmissions in a time division multiple access mode, the "guard time" separating inbound transmissions from interfering with each other can be minimized thereby increasing system efficiency.

In order to accommodate longer message lengths from a remote unit to the master unit, a remote unit can append a request for additional time onto its message to the master unit. The master unit will then compare the priority of the requesting remote unit to the priority of subsequent units and make a decision as to whether to allow the requesting remote unit to use the time division multiple access slots of subsequent units.

By the use of the foregoing, a user can install several host applications employing half duplex polling on a single multidrop network, without a fundamental limit on the number of applications and without the need for extensive engineering adjustment during the installation and maintenance of the system. This allows an end user to have fewer modems or data service units/channel service units (DSU/CSU) on the customer premises.

BRIEF DESCRIPTION OF THE DRAWINGS

FIG. I shows a block diagram of the master unit in the digital application of this invention.

FIG. II shows a block diagram of the remote unit in the digital application of this invention.

FIG. III shows a block diagram of the master unit in the analog application of this invention.

FIG. IV shows a block diagram of the remote unit in the analog application of this invention.

FIG. V shows a schematic of a subframe format.

FIG. VI shows a flow diagram of the initialization of the master unit.

FIG. VII shows a flow diagram of the initialization of the remote unit.

FIG. VIII shows a flow diagram of the normal operation of the master unit.

FIG. IX shows a flow diagram of the normal operation of the remote unit.

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENT

Referring now to the drawings in detail wherein like numerals refer to like elements throughout the several views, master unit 10 for digital applications is shown FIG. 1. Master unit 10 includes a network timing and control processor 12. Network timing and control processor 12 uses firmware or software to implement clock drift reset functions, guard time predict functions, burst receiver control functions, reservation assignment functions, cyclic redundancy check (crc) "checksum" calculations, arq functions, remote transmit control functions and user library update functions such as activity rate and frame parameters. Additionally, network timing

4,937,819

3

and control processor 12 stores user-input initialization parameters including network clock framing periods, slot and subframe assignments, inbound and outbound burst length for each "drop" or remote unit, priority assignments, drop addressing, and port speed assignments.

Reservation request processor 14 allows a drop or remote unit to request more than a single time slot for longer messages. Reservation request processor 14 communicates such a granted request to the network timing and control processor 12.

The aggregate rate multiplexing module 16, in response to commands from the network timing and control processor 12, sets up the timing and bit interleaving of the various application inputs from the various input ports of the primary channel multiplexer 18 and of the overhead and control bits required for outbound control of the remote units from the ranging and network initialization generator 20. The output of the aggregate rate multiplexer 16 is input to the time division multiplexed modulator 22 which transmits the data to the various drops or remote units (see FIG. II). The time division multiplexed modulator 22 is typically a base-band modulator for digital applications.

The master unit 10 also provides, in addition to primary traffic flow over the channel outbound to remote units, a diagnostic channel which can be inband, and any control information necessary to update clock drifts, perform new ranging, etc.

The input in the form of noncontending packets or bursts of information from the various drops or remote units is received via a single data channel by equalizer 24. The input stream has its gain automatically adjusted 26 and is demodulated by demodulator 28.

Demodulator 28 provides digital data to the expansion buffers 30, the ranging receiver 32, the reservation request processor 14 and the bit/ baud timing processor 34.

The expansion buffers 30 output digital data to the data port outputs of the various applications and a diagnostic channel via the port demultiplexer 36.

The ranging receiver 32 receives data from the demodulator 28 during an initial training period (to be described later) so as to store round trip transmission times to each remote unit. This allows an optimization in synchronization of the time division multiple access process, thereby reducing "guard time" between the reception of data from the various remote units thereby increasing the total data transfer rate of the system.

Referring now to FIG. II, data outbound from the time division multiplexed modulator 22 of FIG. I is received by the demodulator 38 of remote unit 40. The output of demodulator 38 is demultiplexed by demultiplexer 42 which, in turn, feeds the primary data receiver ports and any diagnostic or secondary channel required by the particular application. The demodulator 38 also provides bits to the network initialization parameter extraction module 44 and the timing and control block 46. Network initialization parameter extraction module 44 and timing and control block 46 provide the user with network timing extraction information, the slot assignment in which the user is allowed to operate and any control information such as transmit inhibit, transmit enable, reservation grants and other network control provided by the master unit (see FIG. I). Additionally, the ranging calculations, that is, the master unit calculation of the time a signal takes to go from the master unit to any remote unit and vice versa, is stored

4

in the timing and control block 46 where it is used for time advancing the transmit clock.

The timing and control block 46 supplies control bits to compression buffers 48. The timing and control block 46 also feeds the preamble and cyclic redundancy check ("checksum") module 50 and the modulator 52. Additionally, the timing and control block 46 sets up the reservation request generator 54. Reservation request generator 54 monitors the compression buffer for fields exceeding a preset parameter limit which is stored in the initialization parameter table. If a field length exceeding the parameter is sensed, then reservation request generator 54 automatically sets the reservation bits in the outgoing message. The format of the outgoing message, including reservation bits, crc bits, message traffic and preambles are described later herein.

Modulator 52 outputs messages via the communication lines to equalizer 24 of the master unit (see FIG. I).

FIG. III discloses the analog version of the master unit. FIG. I is identical to FIG. III except for the addition of the automatic gain control/equalizer tap store function which is added to ranging receiver 32 so as to provide operating parameters to the burst equalizer 24 and the burst automatic gain control 26 dependent upon the particular remote unit which is transmitting to the master unit. Furthermore, as data is transmitted in "bursts" from the remote unit in the analog mode, such elements as the demodulator, agc, and equalizer in the master unit are entitled "burst demodulators", "burst agc", and "burst equalizer" respectively in the master unit. Similarly, the analog modulator in the remote unit is entitled "burst modulator" (FIG. IV).

This is due to the fact that, in the preferred embodiment, the analog mode requires a true burst mode of operation whereas in the digital case, the burst mode requirement is eliminated due to the presence of a constant envelope baseband which includes idle codes when data is not being transmitted.

It should be noted here that the digital application, for instance, one using Dataphone Digital Service, has an inherent advantage over the analog counterpart, and is therefore a slightly preferred embodiment, in that the inbound traffic received at the master unit is a continuous transmission. That is, either information or idle codes are always being transmitted. There are no periods when the signal is not present. As a result, the number of overhead bits required in each remote transmission in the digital application is considerably reduced. This combined with the high channel to port speed ratio permits lower delay penalties. This results in a greater efficiency in the digital system as compared with the analog system.

The time division multiple access sequence is established by the user. An epoch period or frame is defined by the user. The frame is divided with respect to time into a number of subframes. The subframe is further subdivided into slots, one for each application. Therefore, an application has a preassigned time period (or slot) within a subframe to transmit from the remote unit to the master unit, with the possibility of a reservation request for longer messages.

FIG. V discloses a typical subframe. A subframe is divided into N time slots, each separated by a guard time. Each slot contains a preamble, message bits, reservation request bits, priority bits, and error detection bits. The reservation and priority bits may be replaced by address bits. For example, five bits would permit up to 32 drops. The address bits would be an identifier to

4,937,819

5

the master. The master could then monitor remote clock accuracy, monitor drop transmission events, perform ranging, etc.

Due to the number of total bits exceeding the number of message bits, the aggregate burst (in the case of analog) or transmission (in the case of digital) rate must be higher than the sum of the independent port rates.

As part of the installation of this device, both the master and remote units must be initialized.

As is disclosed in FIG. VI, the first step in the initialization of the master unit is to input the directory of users. This is an address entry for each host application drop.

The next step is to input the address priority level. This is followed by an initialization of such system parameters as frame period, number of time slots per subframe, inbound message lengths, inbound and outbound transmission rates (notice that aggregate burst rate must be higher than the sum of the independent port rates so as to accommodate individual ports along with associated overhead), priority assignments, drop addressing and, port speed assignments. This information is stored in the network timing and control processor.

The initialization phase of operation also includes a ranging calculation for each combination of remote unit (or "drop") and application. The master unit sends a message which makes a round-trip between the master unit and the individual remote unit. The delay period is stored in a library table in the network timing and control processor 12 so that the remote to master unit communication is synchronized, thereby reducing guard time required between successive transmissions and increasing the efficiency of the system.

Additionally, in the case of analog apparatus, the signals received by the master unit during the ranging calculation are used to train the automatic gain control 26 and the equalizer 24, thereby generating an automatic gain control/equalizer tap store function.

Finally, "set-up" parameters are transmitted from the master unit to the remote units so as to establish remote unit transmit clock advances, priority level, and inbound time slot assignments for the primary and diagnostic channels. The remote unit verifies receipt to the master unit.

Initialization of the remote drops is disclosed in FIG. VII. The user inputs the address of the drop or remote unit. The remote unit returns or "loops-back" the ranging signal from the master unit. The remote unit receives the frame period, the number of time slots per subframe, inbound message lengths, port speed and inbound burst (digital) or transmission (analog) rate.

The remote unit receives set-up parameters such as clock advance, priority level and inbound time slot assignment for the primary and diagnostic framing and set-up information to the master unit.

A flow diagram of the normal operation of the master unit is shown in FIG. VIII.

The master unit sends a message to one of the remote units. An inbound message is received from one of the remote units. The address, cyclic redundancy calculation (i.e. checksum) and priority level are checked. If any of these values are invalid, the master unit retransmits to the remote unit. If these values are valid and there is no reservation request, the message is output to the appropriate application host port at the master end. This process is continually repeated for each application in each remote unit.

6

If the aforementioned values are valid but there is a reservation request, then the message is queued in the event that there is a higher priority request or acknowledged and completely received in the event that there is no higher priority request. After the complete message is received, its address, checksum and priority are validated. If these values are valid, the message is output to the host port. If the values are not valid, the initial step of an outbound transmission from the master to the remote is returned.

FIG. IX discloses the normal operation of a drop or remote unit. The remote unit receives and buffers a DTE poll response message. The remote unit computes the reservation request and the CRC (cyclic redundancy check or checksum) bits. The CRC bits provide error detection for both overhead and transmit port primary data bits. The remote unit transmits the first block of data, and possibly a reservation request on the inbound channel at a predetermined time in accordance with the "time advanced" (to allow for transmission time and synchronize so as to reduce guard time) transmit clock. This block of data includes preamble data, reservation request, priority level, delimiter and CRC bits.

From the foregoing, it is seen that this system includes the following features:

1. The master to remote (outbound) transmission is a constant carrier time division multiplexed bit stream in which multiple Host/FEP poll and data traffic is bit interleaved along with master to remote network timing control and diagnostic information.

2. The master unit periodically transmits a network clock reading to all remotes and performs a roundtrip delay transmission calculation ("ranging") to each remote unit. The master unit informs each remote unit of its precise round trip value.

3. The period of the master network clock transmission establishes a "frame". This frame is further segmented into subframes at the remote.

4. The remote establishes a receive clock reference (a delayed version of the Master Network Clock) and a transmit clock reference.

5. For analog applications, the remote sends a long train to the master and the master trains and stores equalization taps and automatic gain control settings unique to the drop. This is unnecessary for digital networks due to the continuous presence of traffic (either idle codes or data).

6. The master unit preassigns time slots within the subframes, one for each of the independent host applications.

7. Upon receiving a poll at the remote DTE, RTS/CTS toggles and the DTE response bits are loaded into a buffer. The remote then transmits these bits in the assigned time slot using the transmit clock reference. Therefore, contention due to inbound poll responses from other remotes is precluded. All inbound transmissions contain a preamble, poll response data bits, reservation request bits, at least one priority bit and error detection bits. In the case of analog networks, the preamble is unique to the remote and enables the master to rapidly set the equalization taps and automatic gain control.

8. All inbound transmissions are at burst rates exceeding the remote port rate.

9. The subframe time slot is sized for the dominant poll response message length for the application. For longer transmissions, the remote unit sets the reserva-

4,937,819

7

tion bits to identify the required number of additional time slots. The priority bit or bits define the remote's relative importance in reducing poll-response delays. The master unit clocks the received message bits to the expansion buffer 30 and checks the reservation and priority bits for error. If no errors are detected, the master responds to the remote with an "authorization to transmit" command and transmits a "transmit inhibit" command to all of the other remote units. Each of these outbound transmissions are error protected so that remote transmission contention is extremely unlikely. The authorized remote commences transmission on the next available time slot and continues until the message transmission has been completed. During this period the expansion buffer clocks data bits to the host.

10. The master unit can recognize whether a remote clock is drifting and so inform the remote with a fast or slow correction value. Such information can be extracted from the actual time of arrival compared with the expected time of arrival at the master unit and the preamble which identifies the transmitted remote.

11. Analysis diagnostic related information is bit interleaved onto the outbound transmission. Because of the relatively slow poll rate of the Analysis system, the inbound response may be assigned to time slot in every Nth subframe. Thus, the aforementioned goals are achieved.

What is claimed is:

1. A communications network comprising:
a master unit;

a plurality of remote units communicating with said master unit in a multidrop configuration;
wherein each of said remote units execute at least one application program, at least one of said remote units executing at least two application programs, said remote units receiving messages outbound from said master unit and responding in a time slot assigned to each of said application programs;

said master unit including a master network timing means with a period which is divided into a plurality of subframes, wherein each subframe is divided into said time slots, and each of said time slots is used as an interval in which one of said application programs in said one of said remote units is assigned to transmit to said master unit in a time division multiple access fashion; and

said master unit including ranging means communicating with said master network timing means wherein a transmission time between said master unit and each of said respective remote units is calculated and transmitted from said master unit to each of said respective remote units, each of said respective remote units using said transmission time to adjust initiation of said time slots.

2. The network of claim 1 wherein said remote units include a reservation request generator which activates a reservation request bit for requesting an additional time interval inbound to said master unit, and wherein said master unit includes a reservation request processor communicating to said master network timing means, said reservation request processor being responsive to said reservation request bit.

3. The network of claim 2 wherein said master unit initiates communication with said remotes using half duplex polled protocols.

4. The network of claim 3 wherein communication from said remote units to said master unit is in the form of modulated bursts on an analog carrier frequency.

5. The network of claim 4 wherein communication from said remote units is received by a burst equalizer in series with a burst automatic gain control in said master unit, said burst equalizer and said burst automatic gain

8

control being responsive to operating parameters which are dependent upon which of said remote units is scheduled to transmit.

6. The network of claim 5 wherein said operating parameters are calculated in an initial training period and stored in said master unit.

7. The network of claim 6 wherein said operating parameters are indexed in said master unit and advanced from said master unit to said burst equalizer and said automatic gain control in response to a preamble unique to each of said remote units which is communicated from said remote units to said master unit.

8. The network of claim 3 wherein communication from said remote units to said master units is in digital form.

9. The network of claim 8 wherein communication between said remote units and said master unit is encoded using baseband modulation.

10. The network of claim 3 wherein said master unit includes a master network clock with a period which is divided into a plurality of subframes, each corresponding to transmission time for one of said remote units, and wherein each subframe is divided into said time slots, and each of said time slots is used as an interval in which one of said applications programs in said one of said remote units transmits to said master unit.

11. The network of claim 2 wherein said time slot comprises a format so as to include a preamble, a poll response data bit, said reservation request bits, at least one priority bit and error detection bit.

12. The network of claim 1 wherein the master unit includes means for calculating clock drifts of the remote units and issuing reset commands to correct the same whereby each remote unit determines its transmit epoch accurately, thereby minimizing guard time while maintaining contention-free transmission to said master unit, said means for calculating clock drifts and issuing reset commands being in communication with said master network timing means.

13. The network of claim 4 wherein said bursts occur at a burst rate which is greater than an aggregate port rate of said remote units.

14. A method for a plurality of remote units to operate a plurality of application programs in communication with a master unit in a multidrop configuration, comprising the steps of:

calculating and storing in said master unit inbound and outbound transmission times between the master unit and said remote units;

dividing a period of a clock in said master unit into a number of subframes, dividing each subframe into a number of slots, each corresponding to transmission times for one of said remote units, and assigning a slot to each of said application programs in said one of said remote units;

transmitting from said master unit to each of said respective remote units the transmission time between said master unit and said respective remote unit, each of said respective remote units using said transmission time to adjust initiation of said slots; and

transmitting data from each of said remote units to said master unit in a time division multiple access configuration wherein each application in each remote unit transmits during said assigned subframe.

15. The method of claim 14 further comprising the step of initiating communication between said master unit and said remote units by using half-duplex polled protocols.

* * * * *